A Shared Control Approach to Safely Limiting Patient Motion Based on Tendon Strain during Robotic-Assisted Shoulder Rehabilitation

Belli, Italo; Hu, Yuxuan; Prendergast, J. Micah; Abbink, David; Seth, Ajay; Peternel, Luka

DOI
10.1109/ICORR66766.2025.11063007

Publication date
2025

Document Version
Final published version

Published in
Proceedings of the International Conference on Rehabilitation Robotics, ICORR 2025

2025 International Conference On Rehabilitation Robotics (ICORR)
May 12-16, 2025. Chicago, IL, USA

A Shared Control Approach to Safely Limiting Patient Motion Based on Tendon Strain During Robotic-Assisted Shoulder Rehabilitation

Italo Belli[1,2,*], Yuxuan Hu[1], J. Micah Prendergast[1], David Abbink[1,3], Ajay Seth[2], and Luka Peternel[1]

***Abstract*— Combining biomechanical modeling with robotic physiotherapy is a promising direction to provide real-time insights during the rehabilitation of patients with musculoskeletal injuries, such as rotator-cuff tears. One aspect is to prevent re-injuries caused by high strain in the injured tissues while allowing patients to perform the required rehabilitation exercises. In this paper, we propose a novel shared control method for robots to limit unsafe patient movements, through physical guidance based on a strain-space representation of the human rotator cuff. The method provides motion corrections through two complementary predictive modules. The first module exerts a lower degree of intervention and is analogous to rumble strips or speed bumps for cars on the road. In this case, an impedance controller induces variable damping to slow down the patient's movement when a danger zone is approached. The second module produces a higher degree of intervention and is analogous to lane-assist in cars. In this case, the robot plans an optimal deflection trajectory and temporarily takes over control of the movement to avoid an unsafe situation. We performed experiments with a healthy participant acting as a patient and evaluated the effect of different human-robot interaction modalities on the resulting human movement in terms of avoidance of high-strain areas of the rotator-cuff tendons and contact forces exchanged.**

I. Introduction

Musculoskeletal injuries can occur during many daily activities, ranging from workplace and household tasks to sports. Such injuries negatively affect the quality of life of those who suffer from them, limiting their comfort, mobility, agency and productivity in their daily life as well as at workplaces. Elderly populations are particularly susceptible to such injuries due to weakening of the musculoskeletal system and general reduction in muscle mass that occurs with age. One of the most typical injuries occurs in rotator cuff muscles that keep the upper arm (humerus) in the socket of the shoulder blade (glenoid of the scapula). The prevalence of rotator cuff injuries is estimated to be as high as 22% in the general population [1].

Traditional therapy in treating injuries to the shoulder involves sessions with human physiotherapists. Nevertheless, re-injuries are common [2], and demand for the therapy often exceeds physiotherapists' capacity, which can further affect the recovery times and completeness of recovery. To this end, robots provide a good solution to augment the capacity of human physiotherapists. Existing works in robotics for physiotherapy enable the generation of prescribed rehabilitation movements [3], as well as learning these movements from expert human demonstrations [4]. These strategies are excellent in automating the movements of human physiotherapists.

Authors are with the Cognitive Robotics[1], BioMechanical Engineering[2] and Sustainable Design Engineering[3] departments at Delft University of Technology, Delft, The Netherlands.
* corresponding author (i.belli@tudelft.nl)

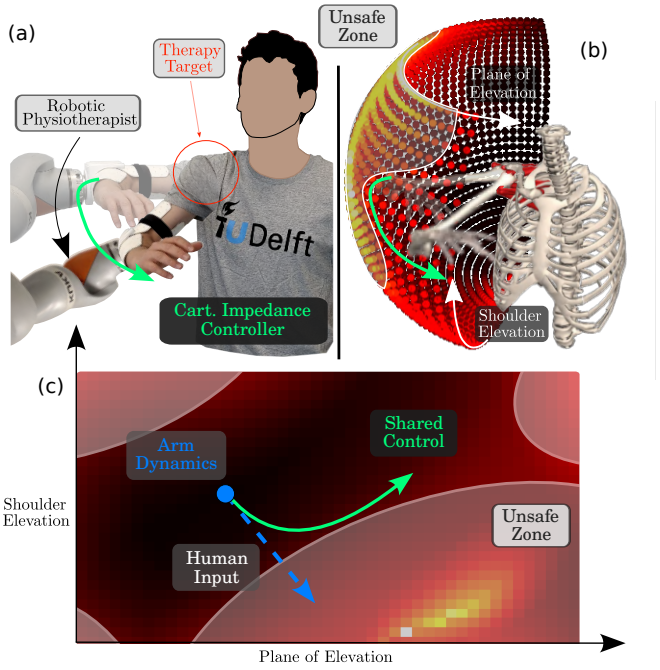

Fig. 1. An overview of the proposed robot-assisted shoulder physiotherapy system. (a) The robot end effector is connected to the human elbow with a customized brace, allowing the robot to limit human motion when necessary. (b) Shoulder strain map is pre-calculated from a high-fidelity musculoskeletal model to determine unsafe movement zones. (c) The current and predicted human arm pose are projected in the strain map to formulate the planning scene. When potential injury is detected, optimal movement in the elbow is achieved through shared control from the robot physiotherapist.

Nevertheless, repeating prescribed or learned movements can be limiting when addressing the specific needs of every patient.

Musculoskeletal models are a valuable addition to the robot motion and force sensors, enabling personalized insights into what is happening inside the human body during the interaction [5]. Most of the existing studies employ models offline, for example, to quantify assistance needed by a human operator [6]–[8] or reduce human metabolic cost in walking-assistive devices [9]–[11]. Some recent studies also explore the use of models in an online manner, where they are integrated into the robot control loops for real-time adjustments of robot assistive movements [12]–[15].

Our previous work [13] developed a method that mapped tendon strains on the human state space using an OpenSim musculoskeletal model to create so-called "strain maps", which chart the high-strain areas where re-injury could occur. The approach employed an offline optimization to create the low-strain therapy trajectories used to navigate the strain maps in real time. A significant advancement was recently made with BATON [16], which enabled an online

2025 International Conference On Rehabilitation Robotics (ICORR) | 979-8-3503-8068-2/25/$31.00 ©2025 IEEE | DOI: 10.1109/ICORR66766.2025.11063007

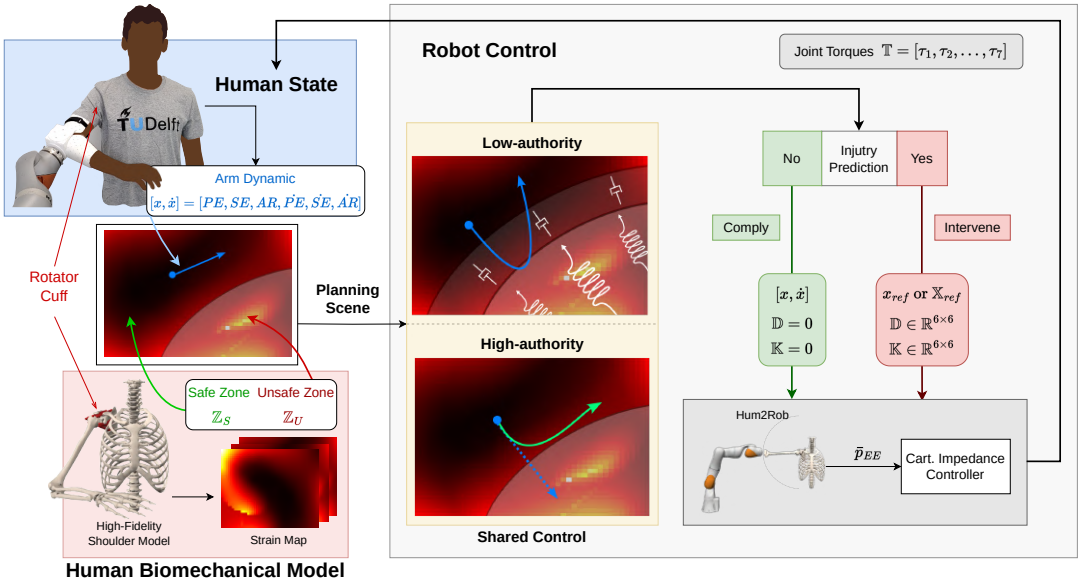

Fig. 2. Diagram of the proposed real-time system for safely limiting human motion during rehabilitation, on the basis of a high-fidelity OpenSim shoulder model. The two shared control modules receive the current state of the human shoulder and their corresponding strain maps to identify dangerous situations and output commands for variable impedance control of the robotic arm. Low-authority guidance is achieved by damping human velocities approaching unsafe high-strain zones. High-authority guidance relies on deflecting human motion away from the danger zones by actively tracking an optimized reference trajectory. These modules are complimentary and can be selected based on the therapy requirements and user preferences.

trajectory optimization and accounted for real-time changes in strain space by considering both the strain behavior and the human arm dynamics. Notably, these approaches focused on optimizing robot-led therapy, which is suitable for early-stage shoulder rehabilitation. Since patients take a more active role as therapy progresses [17], safety and potential re-injury concerns should be addressed while specifically promoting patient-led exercises. In this context, shared control is a promising solution that allows the patient to retain authority over the movement [18], [19], while a robot therapist monitors and potentially intervenes to limit unsafe motions. A simple solution to patient-robot interaction during rotator-cuff rehabilitation was proposed in [15], where the robot provided a "safety net" during physiotherapy by segmenting the dangerous high-strain zones and generating haptic boundaries around them, thus preventing patients or therapists from entering these unsafe regions. Although this offered a promising first step toward a biomechanics-aware safety approach to patient-led physiotherapy, potentially unsafe situations could still occur if a haptic boundary was contacted with high velocity or at a high angle. Therefore, an open challenge is how to predict potential impacts with high-strain zones in advance and avoid them before they occur.

To address this challenge, we developed a method based on shared control that allows the user to perform rehabilitation movements freely, with the robot intervening whenever an incursion into a high-strain danger zone or a high-impact collision with a haptic boundary is about to occur. We propose two distinct predictive modules, each handling the robot's intervention in a different manner. In the first module, the robot exerts lower authority, similar to "rumble strips" or "speed bumps" for cars on the road. In this case, the impedance controller induces variable damping when a danger zone is approached in a potentially unsafe direction to slow down the movement. In contrast, the second module exerts higher authority, akin to lane departure correction autopilot (e.g., "lane assist" or "lane centering") in cars. In this case, the robot plans an optimal deflection trajectory and temporarily takes over control of the movement to avoid an unsafe situation. The two modules are complementary, where the advantage of the first is that it takes away less control from the patient/therapist, while the second module provides a smarter and more assertive way of avoiding unsafe situations. Each module can be selected based on the specific exercise, stage of therapy, user confidence, etc. To demonstrate the key aspects of each, we performed experiments with a healthy participant acting as a patient and evaluated the effect of the different human-robot interactions on the resulting human movement in terms of avoidance of high-strain areas and contact forces exchanged.

II. Methods

Figure 2 illustrates the proposed system with key elements. The human modeling and strain map computation element provide insight into the human musculoskeletal system and allow the physiotherapy robot to navigate away from dangerous high muscle strains. Human state estimation is performed in real-time from the robot's onboard sensors, to monitor human safety. The robot controller implements our two novel modules for predictive guidance to avoid dangerous zones in the strain maps, leveraging a low-level impedance controller that governs physical human-robot interaction. Finally, we present our experimental setup. Each of the key elements is described in the following subsection.

A. Human Biomechanical Model and Strain Maps

To capture the biological properties of the human shoulder, we employ a high-fidelity musculoskeletal model of the

human shoulder complex [20] developed in OpenSim [21], [22]. In particular, we are interested in the behavior of the rotator cuff tendons, spanning the glenohumeral joint and connecting the humerus to the shoulder blade. Therefore, we obtained a reduced-order model of the human shoulder, capturing the mobility of the glenohumeral joint alone in terms of three rotational coordinates (PE: plane of elevation, SE: shoulder elevation, and AR: axial rotation), while retaining the information about the relevant muscles and tendons contained in the original model. As such, the variables that define the configuration of the human model are:

$$\boldsymbol{x} = [\boldsymbol{q},\ \dot{\boldsymbol{q}}]^\top, \quad \text{with} \quad \boldsymbol{q} = [PE, SE, AR]^\top \tag{1}$$

where PE, SE, and AR are defined as the Y-X'-Y'' sequence of intrinsic Euler angles in the fixed shoulder reference frame [20], and $\dot{PE}, \dot{SE}$ and $\dot{AR}$ are their corresponding derivatives with respect to time. Similar to our previous work [13], [16], we employ the OpenSim functionalities to pre-compute the relationship between human pose and strain σ of the tendon(s) of interest, generating "strain maps" that carry the same information as the original model, but much quicker to access. By defining a threshold on the strain value, we leverage this data to identify elliptical unsafe zones $\mathbb{Z}_U$ on the map, corresponding to configurations of the human arm that should be avoided during therapy. Since rehabilitation mostly occurs with a limited speed of movement, the dependency between the fiber velocity and the strain level was neglected.

Movement in the human state space obeys the equations of motion of the corresponding multibody skeletal system. As OpenSim natively provides this information only numerically, we resorted to OpenSimAD [23] as an instrumented version of OpenSim to automatically generate differentiable outcomes and functions from our reduced-order shoulder model. We obtained a differentiable representation of the forward and inverse dynamics functions relating the evolution of the human state $\boldsymbol{x}$ to the generalized torques $\boldsymbol{u} \in \mathbb{R}^3$ applied to the model's degrees of freedom:

$$\dot{\boldsymbol{x}} = f_{\mathrm{FD}}(\boldsymbol{x}, \boldsymbol{u}), \tag{2}$$

$$\boldsymbol{u} = f_{\mathrm{ID}}(\boldsymbol{x}, \dot{\boldsymbol{x}}) \tag{3}$$

The use of model-based human skeletal dynamics was recently proposed as a strategy for real-time planning of biomechanical-aware trajectories on strain maps [16] and brings the benefit of accounting for subject-specific parameters both regarding tendon behavior and human inertial properties.

B. Human State Estimation

We developed a human state estimation module to inform our control algorithms about the current state of the human subject during therapy. Our subject wore a custom-made brace during our experimental validation, connecting them rigidly to the robot's end-effector (Fig. 2). As such, the estimation of the current human state $\hat{\boldsymbol{x}}_t = [\hat{\boldsymbol{q}}_t,\ \hat{\dot{\boldsymbol{q}}}_t]^\top$ can be achieved based on the Cartesian pose $\boldsymbol{p}_{\mathrm{EE}} \in \mathbb{R}^6$ of the end-effector (EE), assuming a fixed orientation of the human torso and negligible movement of the scapula. The interested reader is referred to our previous work for a full derivation of the human pose and velocities ([16], Section II.D).

C. Robot Control

The interaction between our rehabilitation robot and the human subject is shaped with an impedance controller [24] running at 200 Hz, whose parameters are adjusted in real time according to the estimated human state. Specifically, the commanded force at the end-effector $\boldsymbol{F}_{\mathrm{EE}} \in \mathbb{R}^6$ can be defined as:

$$\mathbf{F}_{\mathrm{EE}} = \mathbf{K}_t(\overline{\mathbf{p}}_{\mathrm{EE}} - \mathbf{p}_{\mathrm{EE}}) - \mathbf{D}_t\dot{\mathbf{p}}_{\mathrm{EE}} \tag{4}$$

where the end-effector reference and actual pose are $\overline{\mathbf{p}}_{\mathrm{EE}}$ and $\mathbf{p}_{\mathrm{EE}}$, respectively, and $\mathbf{K}_t,\ \mathbf{D}_t \in \mathbb{R}^{6\times 6}$ are the desired, time-dependent stiffness and damping matrices in Cartesian space. The selection of $\overline{\mathbf{p}}_{\mathrm{EE}}$, $\mathbf{K}_t$, and $\mathbf{D}_t$ ultimately determines the level of authority of the robotic therapist on the movement performed by the human subject, and modifies the control authority of the human during therapy. As such, we present two different modules for tuning the controller's parameters and, effectively, the robot's intervention authority. Both modules combine a prediction of the future human state based on the current human state and determine if future movements will be safe. If the predicted motion will be unsafe (resulting in higher tendon strains), the robot physiotherapy will provide force-based feedback/corrections to the human subject, with lower or higher authority.

1) Low-Authority Shared Control: The human subject retains control of their movements as much as possible, provided that they stay within the safe boundaries defined by the strain maps. In addition to the safe and unsafe zones, we also introduce "damped zones" which represent regions in the human state space that are near, but not yet within, an unsafe zone. At time step t, we perform a biomechanical-safety check on estimated human pose $\hat{\boldsymbol{q}}_t$, projecting it on the strain maps to assess if it lies in a safe, damped, or unsafe zone. Accordingly, we define three conditions (captured in Algorithm 1):

- if $\hat{\boldsymbol{q}}_t$ is safe, then stiffness and damping of the controller are set to zero, permitting free subject movement;
- if the subject's pose lies in a damped zone, then we analyze their velocity $\hat{\dot{\boldsymbol{q}}}_t$. If $\hat{\dot{\boldsymbol{q}}}_t$ points away from the closest unsafe zone, then the movement is safe, and free movement is enabled. Otherwise, the damping in the physical human-robot interaction is increased, setting low stiffness $\mathbf{K}_{\mathrm{low}}$ associated with damping exceeding the critical one $\mathbf{D}_t = 2r\sqrt{\mathbf{K}_{\mathrm{low}}}$ (where we employ the ratio $r > 1$). In this way, the subject receives haptic information for safer direction of movement;
- if $\hat{\boldsymbol{q}}_t$ lies within an unsafe zone, the closest point on the elliptical contour $\mathbb{C}_U$ of the zone is set as a reference point in human coordinates, and tracked with $\mathbf{K}_t = \mathbf{K}_{\mathrm{high}}$, $\mathbf{D}_t = 2\sqrt{\mathbf{K}_{\mathrm{high}}}$. In this way, the human subject is pushed to the nearest safe point by the robot intervention (similar to [15]).

2) High-Authority Shared Control: This modality uses a kinematic prediction of the future human state based on the current human state $\hat{\boldsymbol{x}}_t$, performed over a receding time

Algorithm 1 Low-Authority Shared Control

Given: Unsafe Zone $\mathbb{Z}_U$, Damped Zone $\mathbb{Z}_D$
Input: Human State $\boldsymbol{x}_t = [\boldsymbol{q}_t, \dot{\boldsymbol{q}}_t]$
if $\boldsymbol{q}_t \in \mathbb{Z}_D \setminus \mathbb{Z}_U$ **AND** $\dot{\boldsymbol{q}} \to \mathbb{Z}_U$ **then**
  $\boldsymbol{K}_t \leftarrow \boldsymbol{K}_{low}$
  $\boldsymbol{D}_t \leftarrow 2r\sqrt{\boldsymbol{K}_{low}}$
  $\boldsymbol{x}_{\text{ref}} = \boldsymbol{x}_t$
else if $\boldsymbol{q} \in \mathbb{Z}_U$ **then**
  $\boldsymbol{K}_t \leftarrow \boldsymbol{K}_{high}$
  $\boldsymbol{D}_t \leftarrow 2\sqrt{\boldsymbol{K}_{high}}$
  $\boldsymbol{x}_{\text{ref}} = \boldsymbol{x}^\star$, with $\boldsymbol{x}^\star = \textbf{argmin}_{\boldsymbol{x}}$ distance$(\boldsymbol{x}_t, \mathbb{C}_U)$
else
  $\boldsymbol{K}_t \leftarrow \boldsymbol{0}$
  $\boldsymbol{D}_t \leftarrow \boldsymbol{0}$
end if
Return: $[K,\ D,\ \boldsymbol{p}_{\text{EE}}(\boldsymbol{x}_{\text{ref}})]$ to Robot Control

horizon T, to find the estimated future subject's trajectory $\hat{\boldsymbol{X}} = [\hat{\boldsymbol{x}}_{t+1},\ \hat{\boldsymbol{x}}_{t+2}, ...,\ \hat{\boldsymbol{x}}_{t+T}]^\top$. The future trajectory is analyzed to determine whether it intersects any unsafe zone. When $\hat{\boldsymbol{X}}$ is entirely safe, free movement is enabled by setting zero Cartesian stiffness and damping matrices, similarly to II-C.1. If, instead, the current movement leads the subject to any unsafe zones, the algorithm determines an alternative trajectory $\boldsymbol{X}^\star$ by defining an optimization problem that minimally deflects the predicted trajectory but ensures that no unsafe configuration is reached. The robot temporarily takes control of the motion, setting $\mathbf{K}_t = \mathbf{K}_{\text{high}}$, $\mathbf{D}_t = 2\sqrt{\mathbf{K}_{\text{high}}}$, and tracks $\boldsymbol{X}^\star$ so the human is deflected to a safe configuration where free movement is enabled again (see Algorithm 2).

Below, we specify the cost function and constraints that constitute the optimal control problem to be solved over the generic time interval $[t, t+T]$ to find $\boldsymbol{X}^\star$.

Cost Function: we propose three terms that an optimal deflection of the human movement should minimize, ensuring that the latest human intention is respected as much as possible and that forces and accelerations be low to mitigate human discomfort:

- $L_1(\boldsymbol{x}_t, \hat{\boldsymbol{x}}_t) = w_{\text{pos}}\gamma^t||\hat{\boldsymbol{q}}_t - \boldsymbol{q}_t||_2^2 + w_{\text{vel}}||\hat{\dot{\boldsymbol{q}}}_t - \dot{\boldsymbol{q}}_t||_2^2$, which weights the distance between a given human state and the predicted one, potentially with different weights for the human coordinates position and velocities. Note that a discount factor $\gamma < 1$ is employed, to permit larger deviations as time progresses;
- $L_2(\hat{\boldsymbol{x}}_t, \boldsymbol{u}_t) = w_{\text{torq}}||\boldsymbol{u} - f_{\text{ID}(\hat{\boldsymbol{x}}_t, \boldsymbol{0})}||_2^2$, accounting for the difference between the generalized torques exerted on the human model and the torques that would guarantee equilibrium of the future trajectory;
- $L_3(\dot{\boldsymbol{x}}_t) = w_{\text{acc}}||\dot{\boldsymbol{q}}_t||_2^2$, weighting the instantaneous acceleration of the human model.

Constraints: the set of constraints that we employ formalize the requirements that $\boldsymbol{X}^\star$ should satisfy:

- *Initial Condition*: Trivially, the deflected trajectory should start from the current human state;
- *Tendon Safety*: The human pose should not cause excessive strain on the healing tendon(s), leading to the requirement $\boldsymbol{q}_t \notin \mathbb{Z}_U$;
- *Dynamic consistency*: For each instant in time, the system's accelerations need to respect the dynamics of the human skeletal system, resulting in $\dot{\boldsymbol{x}}_t = f_{\text{FD}}(\boldsymbol{x}_t, \boldsymbol{u}_t)$;
- *Terminal Conditions*: At the end of the optimization interval, we require that the human speed be zero and that the torques on the human model match the equilibrium ones. This results in $|\dot{\boldsymbol{q}}_T| \leq \varepsilon_{\text{vel}}$, $||\boldsymbol{u}_T - f_{\text{ID}}(\boldsymbol{x}_T, \boldsymbol{0})||_2^2 \leq \varepsilon_{\text{torq}}$, where suitable tolerances $\varepsilon_{\text{vel}}, \varepsilon_{\text{torq}}$ are used. Moreover, the final human state should be far enough from the contour of the closer unsafe zone, to guarantee that the movement can be continued safely by the human subject: dist$(\boldsymbol{x}_T, \mathbb{C}_U) \geq \delta$.

Through orthogonal collocation techniques [25], the problem formalized above is cast into an equivalent Non-Linear Programming problem (NLP) that structure-exploiting solvers can solve.

For slow movements and relatively short T, the human skeletal dynamics specified by (2)-(3) could introduce more involved computations without significantly affecting the quality of the optimal deflection $\boldsymbol{X}^\star$. We tested this hypothesis by considering an alternative version of the NLP where the strain maps are navigated by a virtual point-mass system instead, actuated by ideal forces. The accuracy of this approximation, evaluated through computer simulations, is presented in Section III-A.

Algorithm 2 High-Authority Shared Control

Given: Unsafe Zone $\mathbb{Z}_U$
Input: Human State $\boldsymbol{x}_t$
$\hat{\boldsymbol{X}} \leftarrow$ **pathPrediction**$(\boldsymbol{x}_t)$
if $\exists \boldsymbol{x}_i \in \hat{\boldsymbol{X}} \mid \boldsymbol{x}_i \in \mathbb{Z}_U$ **then**
  $\boldsymbol{K}_t \leftarrow \boldsymbol{K}_{high}$
  $\boldsymbol{D}_t \leftarrow 2\sqrt{\boldsymbol{K}_{high}}$
  $\boldsymbol{X}^\star_{ref} \leftarrow \textbf{argmin}_{\boldsymbol{X}_{\text{ref}}} \mathcal{L}(\boldsymbol{X}, \boldsymbol{X}_{\text{ref}})$
else
  $\boldsymbol{K}_t \leftarrow \boldsymbol{0}$
  $\boldsymbol{D}_t \leftarrow \boldsymbol{0}$
end if
Return: $[K,\ D,\ \boldsymbol{p}_{\text{EE}}(\boldsymbol{X}^\star_{\text{ref}})]$ to Robot Control

D. Initial User Testing

Before conducting proof-of-concept experiments, we solicited the input of a physiotherapist for qualitative feedback on the interaction. Initial trials with the PT were deemed to be overly "aggressive" during motion correction. To account for this, steps were taken to better tune the controller parameters, (i.e., tuning the weights in the cost function and increasing the sampling frequency of the reference trajectory) to avoid this perceived behavior.

E. Experimental setup

A healthy subject participated in our experiments and interacted with a KUKA LBR iiwa 7 through our custom arm brace, allowing simultaneous human state estimation and physical human-robot collaboration during a simulated rehabilitation session. Our experimental protocol was approved by the Human Research Ethics Committee of TU

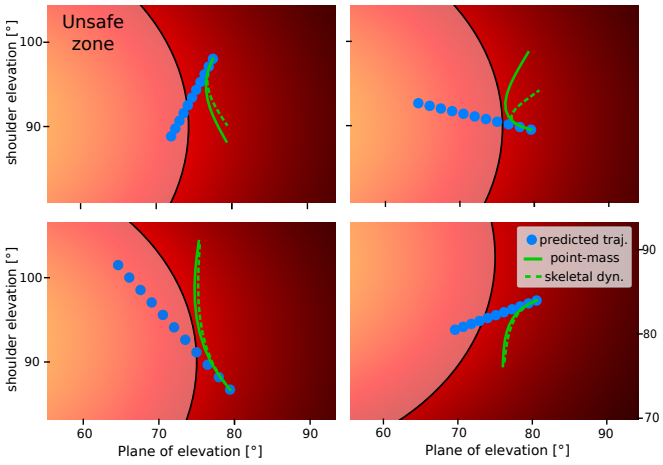

Fig. 3. Simulation comparison of optimal deflections generated by the high-authority module, contrasting the results when using the human model's dynamics (dashed green) versus a virtual point mass (solid green) for navigating the strain map.

Delft. To demonstrate the key functionalities of our method and ensure the repeatability of the experiments, we locked the AR degree of freedom and considered a custom strain map with one time-invariant unsafe zone. The subject was instructed to maintain a constant torso orientation during the interaction to satisfy the assumption of Section II-B. We selected $k_{high,t} = 400\frac{N}{m}$ and $k_{high,r} = 20\frac{N}{[rad]}$ as high translational and rotational stiffness respectively, while low values were set to $k_{\mathrm{low,t}} = 20\frac{N}{m}$ and $k_{\mathrm{low,r}} = 5\frac{N}{[rad]}$.

By removing the damped zone in the low-authority module, we effectively recreated the haptic boundaries considered in our previous work [15], to measure the performances of this baseline controller in terms of robot-commanded forces when the unsafe zone was reached (Section III-B). Then, we set $r = 4$ to increase the damping of the controller, so that our subject could receive anticipatory haptic feedback before impacting the unsafe zone. In the high-authority module, we selected $T = 1$ s and divided it into 10 discrete steps, and up-sampled the resulting optimal deflection to the required frequency for the controller. When the deflection was executed, the system also played a sound to promptly inform the subject that they should comply with the robot's corrective movement.

During these proof-of-concept experiments, the subject interacted with the robot by moving their arm through a natural range of motion. They were then directed to move towards the prescribed unsafe region while the resulting trajectory and forces were recorded. The subject repeated this for both deflection modalities to analyze the effectiveness of our method. Computations were performed on a Dell Latitude 7420 laptop with an i7-1185G7 processor, interfaced with a Dell workstation with an Xeon W-2123 processor dedicated to the impedance controller. Our code is available at *https://github.com/itbellix/biomechanical_safe_deflection*.

III. Results

A. Evaluation of the approximated dynamics

We employed computer simulations to evaluate the differences between the optimal deflections produced by the NLP presented in Section II-C.2 when the strain maps are navigated with the human model's dynamics or by the virtual point mass. We selected 100 random initial conditions for human poses and velocities, to represent the different ways that a subject would move toward an unsafe zone, and computed the optimal deflection proposed by the NLP in the two cases. Figure 3 exemplifies the results obtained, given the kinematic prediction of the human trajectory computed from the randomly selected initial state. Overall, we observed a RMSE of about $2°$ between the last point of the two trajectories. The NLPs employing human skeletal dynamics required an average solution time of 740 ms (failing to converge in $4/100$ instances), whereas the NLPs with virtual point mass approximation could be solved in 12 ms on average, with 100% convergence rate. For these reasons, we opted to use the second approximation in the experiments performed with our physical robot, the results of which we present next.

B. Baseline controller

We ran our baseline controller [15] and recorded the forces generated by the robot on the human as a consequence of collisions with the unsafe zone during subject-led movements. The peak force magnitude that we observed exceeded 17 N, with the subject executing similar movements as during the rest of our experiments, but receiving no feedback or deflection from the robot before hitting the unsafe zone.

C. Shared control modalities

We analyzed our two shared controllers during physical human-robot interaction, mimicking a physiotherapy session. In both cases, the human subject moved their arm, and the robot behavior was dictated alternatively by one of the two algorithms presented in Sections II-C.1, II-C.2. As the movement started from a safe region, initially both modules guaranteed minimal interaction force. On the contrary, the behavior was different as the unsafe zone was approached. The lower authority module (Figure 4) damped the human movement as the subject entered the damped zone, generating a dragging force opposing human velocities pointing towards the unsafe zone. The subject continued their movement and entered the unsafe zone, triggering the algorithm to pull them toward the closest point on the zone contour. Finally, the subject left the unsafe zone and moved away from it, receiving no further robot intervention in virtue of the safe direction of movement.

Regarding the higher authority module (Figure 5), the subject could move undisturbed closer to the unsafe zone until their predicted future trajectory was no longer safe. At this stage, the algorithm proposed a minimal deflection of the original trajectory, triggering the robot to lead the human subject toward a position where they could move safely again. For both cases, we analyzed the magnitude of the Cartesian forces that were commanded to the human subject and found that the first control method led to a peak interaction force of about 8 N, while a maximum force of about 6 N was reached with the second control method.

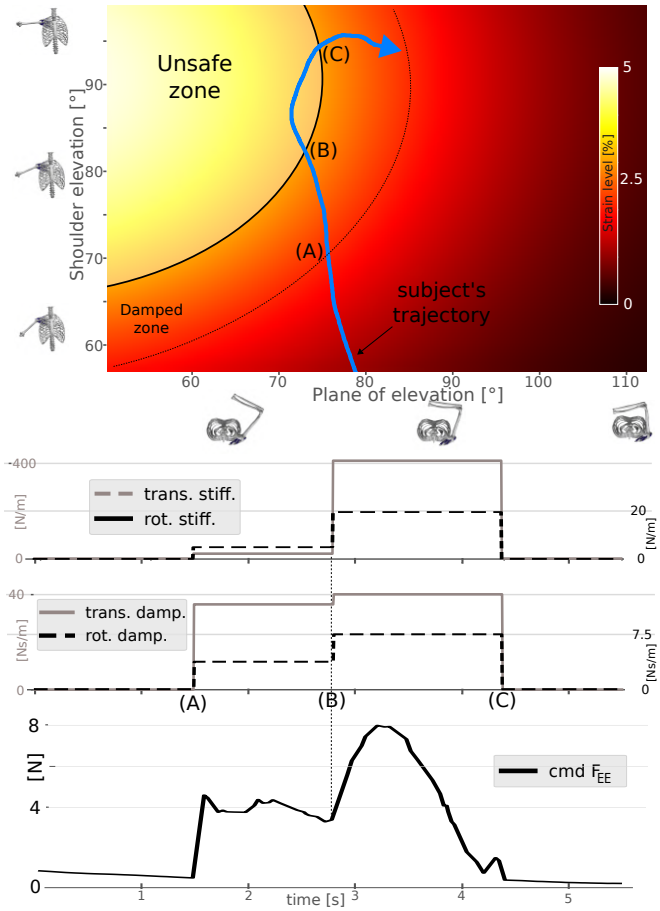

Fig. 4. Low-authority shared control module behavior, during physical human-robot interaction. Top: strain map in which the subject moved, with skeletal poses detailing the axes. The subject approached the unsafe zone: (A) first, the "damped zone" was encountered, then (B) they entered the "unsafe zone" and experienced higher robot correction before leaving the zone and continuing with their free movement (C). Evolution of the controller's stiffness and damping and resulting interaction force are highlighted.

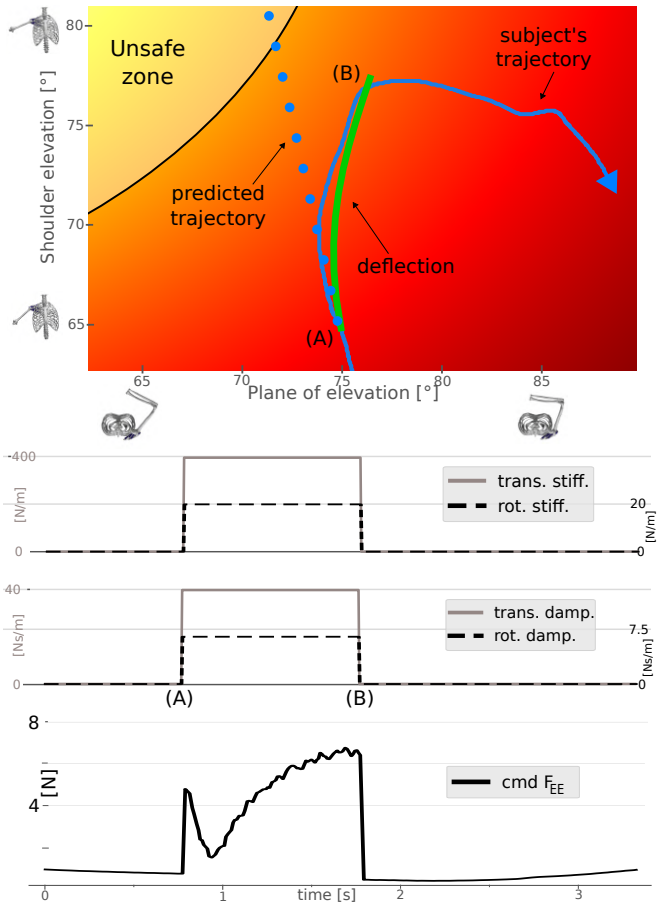

Fig. 5. High-authority shared control module behavior, during physical human-robot interaction. Top: strain map in which the subject moved, with skeletal poses detailing the axes. The subject approached the unsafe zone, triggering the system to impose a short deflection on their movement (A). After that (B), control was given back to the human, who could independently continue their free motion. The evolution of the controller's stiffness and damping and the resulting interaction force are highlighted.

IV. Discussion and Conclusions

We have presented two shared control modalities that allow a robot to limit unsafe human movement during robotic-assisted shoulder rehabilitation, using information from a high-fidelity biomechanical model. Both modules permit human-led movements to maximize patients' independent exploration of their shoulder range of motion, safeguarding their users against reaching shoulder configurations that could generate unsafe levels of strain in injured/healing tendons. When the current human action is deemed unsafe, our controller intervenes with different degrees of authority to shape the resulting human movement. The use of the shared control paradigm expands our previous contributions, in which the robotic therapist unilaterally dictated the therapeutic motion [16] and did not anticipate human intentions [15].

In the present work, the low-authority control modality, designed to allow the human subject to retain greater agency over the resulting movement, was capable of providing state-dependent haptic feedback during human motion. The system produced a force field that can allow the user to avoid unsafe zones autonomously (Figure 4). While entering such zones is still possible, the maximum robot-commanded force was smaller than what was observed when no predictive damping was provided to the user for similar movements, since higher unsafe velocities were limited. The high-authority control modality, on the other hand, is expected to reduce the autonomy of the human subject, as it temporarily takes over the movement, to reposition the human in a pose where they can safely continue their therapy. This approach allowed the system to avoid our experimental unsafe zone better, with lower force produced by the robot, by guaranteeing initial alignment between the human intention and the robot's deflection. Moreover, our real-robot results validated the initial simulation findings reported in Figure 3, showing that the subject can be smoothly deflected without accounting for the specific human skeletal dynamics, for the type of rehabilitation movements considered.

Simplifying human inertia also allows for reduced computational costs. Comparing the optimal deflections obtained with and without accounting for the human skeletal dynamics highlighted minor differences in our simulation results (see Figure 3), likely due to different relative weighting of

the torque terms in the cost function. Indeed, the inertial parameters of virtual point mass navigating across the strain map were not tuned to our participant since only the human kinematics was employed by the robot's impedance controller. Identifying the human-arm inertia at the robot's end-effector and using these parameters in the point mass approximation could enhance the alignment between the two methods under low accelerations, as in our case. This approach could also enable integrating optimized generalized torques from the human model into the robot controller, such as for personalized gravity compensation [16].

While our findings demonstrate the potential of the proposed shared control modalities, there are limitations that warrant consideration and open avenues for interesting research. In this work, we considered a stationary exemplary strain map, purely dependent on the human shoulder pose, which simplified the demonstration of our control modalities. However, the effects of tendon fiber velocity and muscle activation on the resulting strain are currently ignored, and their inclusion will be necessary in the future, as was demonstrated previously [26]. Including rapid estimation of the muscle activation through dedicated solvers such as [27] is an interesting future direction, as it would improve the quality of the strain estimations by accounting directly for the effect of the interaction wrenches that are exchanged between the human and the robot during their physical interaction.

We are also planning to further explore the control modality that potential users/stakeholders (both patients and therapists) might prefer, allowing us to better define the most effective way to share control of therapy between human and robotic agents. As noted in the methodology, to gather feedback on the system during development, we have conducted a preliminary interactive session with one of the leading experts in orthopedic physiotherapy in the Netherlands. The response was generally very positive, and the main takeaway was to make the shared control system slightly less aggressive. We accounted for this feedback prior to doing experiments for this paper, but our future work will need to focus on testing the method on multiple healthy participants first, to further explore the user experience before moving to tests in more clinical settings.

ACKNOWLEDGMENTS

We thank Michael J. Davidson for the useful discussions about the PTbot project at TU Delft. We acknowledge the support from the transdisciplinary research and innovation centre FRAIM and the Chan Zuckerberg Initiative DAF, an advised fund of Silicon Valley Community Foundation through grants 2020-218896 and 2022-252796.

REFERENCES

[1] H. Minagawa, N. Yamamoto, H. Abe, M. Fukuda, N. Seki, K. Kikuchi, H. Kijima, and E. Itoi, "Prevalence of symptomatic and asymptomatic rotator cuff tears in the general population: from mass-screening in one village," *J. Orthop.*, vol. 10, no. 1, 2013.

[2] F. V. Sciarretta, D. Moya, and K. List, "Current trends in rehabilitation of rotator cuff injuries," *Sicot-j*, vol. 9, 2023.

[3] Y. Cen, J. Yuan, S. Ma, J. Luo, and H. Wang, "Trajectory optimization algorithm of trajectory rehabilitation training mode for rehabilitation robot," in *IEEE ROBIO)*, IEEE, 2022.

[4] B. Luciani, L. Roveda, F. Braghin, A. Pedrocchi, and M. Gandolla, "Trajectory learning by therapists' demonstrations for an upper limb rehabilitation exoskeleton," *IEEE RA-L*, 2023.

[5] C. Fang, L. Peternel, A. Seth, M. Sartori, K. Mombaur, and E. Yoshida, "Human modeling in physical human-robot interaction: A brief survey," *IEEE RA-L*, 2023.

[6] M. G. Carmichael and D. Liu, "Estimating physical assistance need using a musculoskeletal model," *IEEE TBME*, vol. 60, no. 7, 2013.

[7] B. Ghannadi, N. Mehrabi, R. S. Razavian, and J. McPhee, "Nonlinear model predictive control of an upper extremity rehabilitation robot using a two-dimensional human-robot interaction model," in *IEEE/RSJ IROS*, IEEE, 2017.

[8] A. Zignoli, F. Biral, K. Yokoyama, and T. Shimono, "Including a musculoskeletal model in the control loop of an assistive robot for the design of optimal target forces," in *IEEE IECON*, vol. 1, IEEE, 2019.

[9] J. Fang and Y. Yuan, "Human-in-the-loop optimization of wearable robots to reduce the human metabolic energy cost in physical movements," *Robot. Auton. Syst.*, vol. 127, 2020.

[10] D. F. Gordon, C. McGreavy, A. Christou, and S. Vijayakumar, "Human-in-the-loop optimization of exoskeleton assistance via online simulation of metabolic cost," *IEEE T-RO*, vol. 38, no. 3, 2022.

[11] K. Li, M. Tucker, R. Gehlhar, Y. Yue, and A. D. Ames, "Natural multicontact walking for robotic assistive devices via musculoskeletal models and hybrid zero dynamics," *IEEE RA-L*, vol. 7, no. 2, 2022.

[12] P. K. Jamwal, S. Hussain, Y. H. Tsoi, and S. Q. Xie, "Musculoskeletal model for path generation and modification of an ankle rehabilitation robot," *IEEE Trans. Hum.-Mach. Syst.*, vol. 50, no. 5, 2020.

[13] J. M. Prendergast, S. Balvert, T. Driessen, A. Seth, and L. Peternel, "Biomechanics aware collaborative robot system for delivery of safe physical therapy in shoulder rehabilitation," *IEEE RA-L*, vol. 6, no. 4, 2021.

[14] G. Clark and H. B. Amor, "Learning ergonomic control in human–robot symbiotic walking," *IEEE T-RO*, vol. 39, no. 1, 2022.

[15] S. Balvert, J. M. Prendergast, I. Belli, A. Seth, and L. Peternel, "Enabling patient-and teleoperator-led robotic physiotherapy via strain map segmentation and shared-authority," in *IEEE-RAS Humanoids*, IEEE, 2022.

[16] I. Belli, J. M. Prendergast, A. Seth, and L. Peternel, "Biomechanics-aware trajectory optimization for navigation during robotic physiotherapy," *arXiv preprint*, 2024.

[17] T. A. Sgroi and M. Cilenti, "Rotator cuff repair: post-operative rehabilitation concepts," *Curr Rev Musculoskelet Med*, vol. 11, 2018.

[18] D. A. Abbink, T. Carlson, M. Mulder, J. C. De Winter, F. Aminravan, T. L. Gibo, and E. R. Boer, "A topology of shared control systems—finding common ground in diversity," *IEEE Trans. Hum.-Mach. Syst.*, vol. 48, no. 5, 2018.

[19] D. P. Losey, C. G. McDonald, E. Battaglia, and M. K. O'Malley, "A review of intent detection, arbitration, and communication aspects of shared control for physical human–robot interaction," *Applied Mechanics Reviews*, vol. 70, no. 1, 2018.

[20] A. Seth, M. Dong, R. Matias, and S. Delp, "Muscle contributions to upper-extremity movement and work from a musculoskeletal model of the human shoulder," *Front. neurorob.*, vol. 13, 2019.

[21] S. L. Delp, F. C. Anderson, A. S. Arnold, P. Loan, A. Habib, C. T. John, E. Guendelman, and D. G. Thelen, "Opensim: open-source software to create and analyze dynamic simulations of movement," *IEEE TBME*, vol. 54, no. 11, 2007.

[22] A. Seth, J. L. Hicks, T. K. Uchida, A. Habib, C. L. Dembia, J. J. Dunne, C. F. Ong, M. S. DeMers, A. Rajagopal, M. Millard, *et al.*, "Opensim: Simulating musculoskeletal dynamics and neuromuscular control to study human and animal movement," *PLoS Comput. Biol.*, vol. 14, no. 7, 2018.

[23] A. Falisse, G. Serrancolí, C. L. Dembia, J. Gillis, and F. De Groote, "Algorithmic differentiation improves the computational efficiency of opensim-based trajectory optimization of human movement," *PLoS One*, vol. 14, no. 10, 2019.

[24] N. Hogan, "Impedance control: An approach to manipulation," in *American control conference*, IEEE, 1984.

[25] D. Garg, M. Patterson, W. Hager, A. Rao, D. R. Benson, and G. T. Huntington, "An overview of three pseudospectral methods for the numerical solution of optimal control problems," *hal-01615132*, 2017.

[26] I. Beck, I. Belli, L. Peternel, A. Seth, and J. M. Prendergast, "Real-time tendon strain estimation of rotator-cuff muscles during robotic-assisted rehabilitation," in *IEEE Humanoids*, IEEE, 2023.

[27] I. Belli, S. Joshi, J. M. Prendergast, I. Beck, C. Della Santina, L. Peternel, and A. Seth, "Does enforcing glenohumeral joint stability matter? a new rapid muscle redundancy solver highlights the importance of non-superficial shoulder muscles," *PLoS one*, vol. 18, no. 11, 2023.