Biomechanics-aware control for robot-assisted physiotherapy

A novel approach to treating shoulder injuries

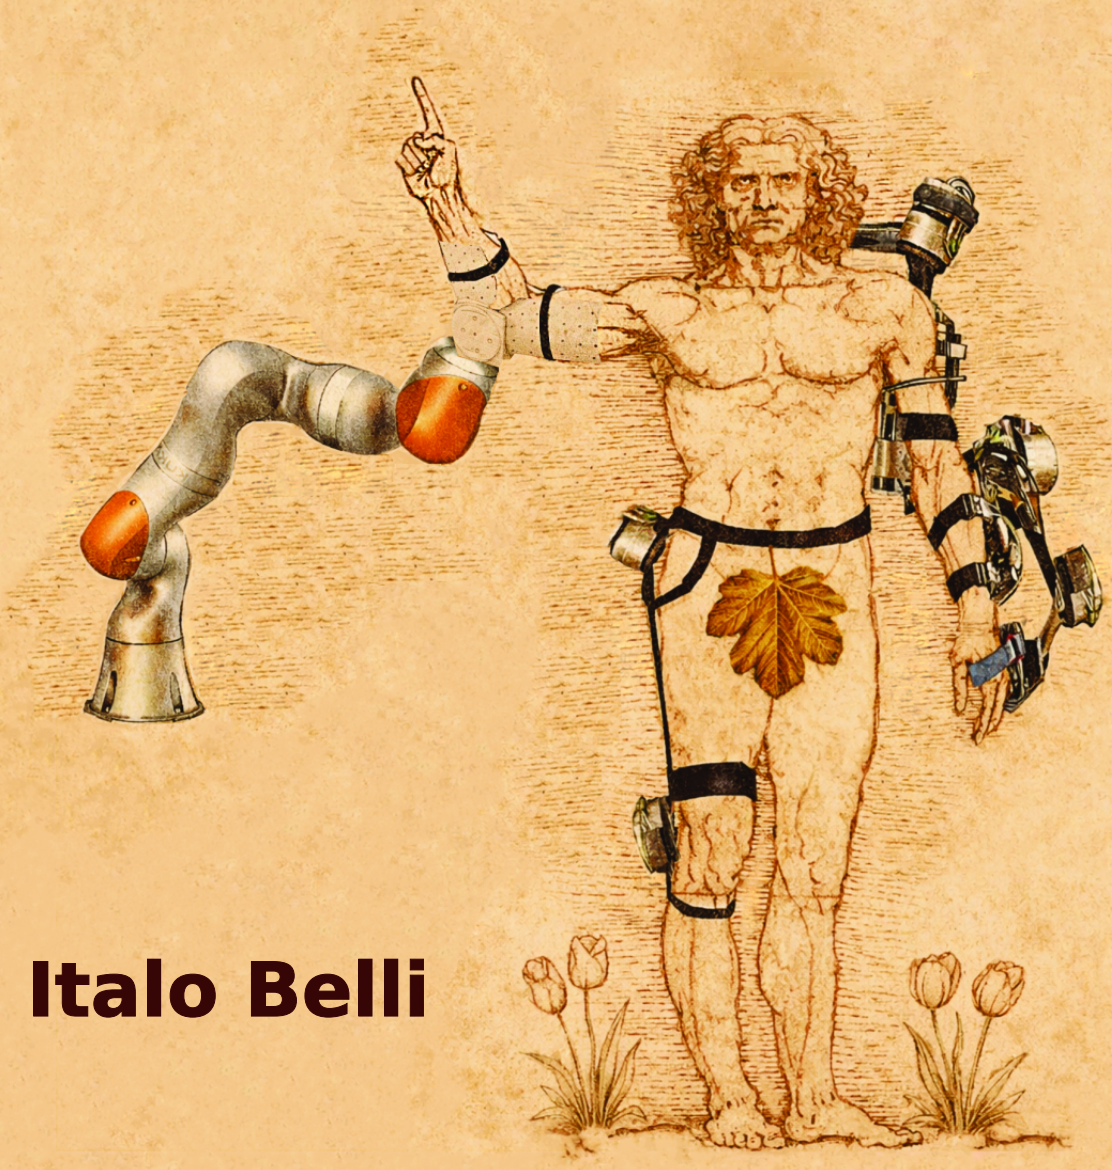

Italo Belli

Propositions

accompanying the dissertation

Biomechanics-aware control for robot-assisted physiotherapy

A novel approach to treating shoulder injuries

by

Italo BELLI

1. For rehabilitation robots to surpass the capabilities of human physiotherapists, their control algorithms need to consider human biomechanics. (This thesis)
2. Non-invasive monitoring of musculoskeletal metrics in physical human-robot interaction is only feasible through model-based techniques. (This thesis)
3. Assistive robots' abilities to safely react to human movement are far more important than their abilities to predict it. (Chapter 3)
4. Musculoskeletal simulations, performed without sensitivity analysis of either modeling or experimental parameters, cannot deliver new scientific insights. (Chapter 2)
5. Replacing human labor with robots can not be justified solely by arguments around reduced physical effort and repetitive tasks.
6. The most effective way to improve research software performance during a PhD is by upgrading hardware and using commercial software.
7. Interdisciplinary research is harder to advance, since it must meet the demands of several fields at once while belonging fully to none.
8. In the absence of unlimited resources, academic research inevitably reflects the priorities of those with power, proximity, or capital.
9. The only way to eradicate systemic discrimination is to favour underrepresented communities, in academia as well as in our societies.
10. Life should not be put on hold because of a PhD.

These propositions are regarded as opposable and defendable, and have been approved as such by the promotor Prof. dr. ir. D. Abbink and by the copromotors Dr. ir. L. Peternel and Dr. ir. A. Seth.

Biomechanics-aware control for robot-assisted physiotherapy

A novel approach to treating shoulder injuries

Biomechanics-aware control for robot-assisted physiotherapy

A novel approach to treating shoulder injuries

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology
by the authority of the Rector Magnificus, Prof. dr. ir. H. Bijl,
chair of the Board for Doctorates
to be defended publicly on
Friday 13 February 2026 at 10:00

by

Italo BELLI

This dissertation has been approved by the promotors.

Composition of the doctoral committee:

Rector magnificus,	chairperson
Prof. dr. ir. D. A. Abbink,	Delft University of Technology, *promotor*
Dr. L. Peternel,	Delft University of Technology, *copromotor*
Dr. ir. A. Seth,	Delft University of Technology, *copromotor*
Independent members:	
Prof. dr. F.C.T. van der Helm	Delft University of Technology
Prof. dr. ir. M. Sartori	University of Twente
Dr. ing. L. Marchal Crespo	Delft University of Technology
Dr. P. Wolf	ETH Zürich, Switzerland
Prof. dr. ir. M. Wisse,	Delft University of Technology, *reserve member*

Prof. dr. ir. J.M. Prendergast heeft als begeleider in belangrijke mate aan de totstandkoming van het proefschrift bijgedragen.

This research was partially funded by the Chan Zuckerberg Initiative DAF, an advised fund of Silicon Valley Community Foundation, through grants 2020-218896 and 2022-252796, and received support from the transdisciplinary research and innovation centre FRAIM. The research visit at the Sensory-Motor Systems Lab, ETH Zürich, was supported by the IDEA League Student Research Grant.

Keywords: Rehabilitation Robotics, Musculoskeletal Modeling, Physical Human-Robot Interaction, Optimization, Optimal Control, Muscle Redundancy, Rotator Cuff, Shoulder Physiotherapy.

Printed by: Gildeprint – www.gildeprint.nl

Cover by: Italo Belli and Maia Rigot, with the use of generative AI (ChatGPT).

ISBN (paperback) 978-94-6384-906-7
ISBN (e-book) 978-94-6518-244-5

An electronic copy of this dissertation is available at
https://repository.tudelft.nl/.

“Thy firmness makes my circle just,
And makes me end where I begun.”

Contents

Summary

Musculoskeletal injuries are among the leading causes of pain, disability, and loss of independence worldwide. They affect millions of people, with prevalence rising steeply with age. One of the most common musculoskeletal injuries is tears to the shoulder rotator cuff. As these muscle-tendon tissues are anatomically constricted in a very narrow space between the shoulder bones, they are frequently subject to trauma or wear. Treatment of these injuries is both medically and socially pressing: they impair daily activities, limit the ability to work and engage in sports, and generate high personal and healthcare costs. Rehabilitation is essential to recovery, but it is often lengthy and labor-intensive for both physiotherapists (PTs) and patients. Moreover, it is prone to setbacks such as re-injury, since PTs lack quantitative tools to monitor the evolution of complex musculoskeletal structures during therapy.

In this context, the adoption of robotic devices offers opportunities to support manual manipulation of patients and provide sophisticated sensors to monitor them. Yet, despite advances in robot design and control, current systems remain unaware of the patient's underlying biomechanics, and therefore cannot monitor or prevent harmful loading of healing tissues.

This thesis addresses such critical lack of knowledge by embedding state-of-the-art musculoskeletal models into the control of rehabilitation robots. Through the development of novel algorithms, it enables real-time estimation of deep muscle activity and tendon strain in the shoulder during physical human-robot interaction. By spanning from improved biomechanical simulations to their integration in robotic therapy execution, this work significantly advances the current state of the art to form a cohesive framework for **biomechanics-aware robotic physiotherapy**.

First, Chapter 1 situates the work within the current scientific and clinical landscape, highlighting the limitations of existing rehabilitation approaches and motivating the integration of biomechanics into robotics.

Chapter 2 introduces the Rapid Muscle Redundancy (RMR) solver, a novel algorithm to rapidly estimate physiologically consistent muscle activations in the shoulder through a computational model. By enforcing glenohumeral stability and accounting for passive fiber contributions, the solver demonstrated that rotator cuff activity cannot be inferred from superficial muscles alone, underscoring the need for simulations to monitor deep muscle engagement.

Chapter 3 builds on this foundation with BATON (Biomechanics-Aware Trajectory Optimization for Navigation of tendon strain maps), which integrates the RMR solver into robotic control. BATON enables online estimation of tendon strain during human–robot interaction, generating safe and adaptive rehabilitation trajectories that reduce re-injury risk.

Chapter 4 elaborates further on our approach to estimating tendon strain maps,

which capture how kinematics, external forces, and activations jointly affect tendon strain. We present how these maps can estimate rapidly the effect of robotic resistance to human movement by visualizing the effect of exercises, offering a new diagnostic and personalization tool to PTs as well.

Chapter 5 addresses rehabilitation stages where patients are encouraged to move actively but remain vulnerable to re-injury. It introduces shared-control strategies where the robot allows free motion in safe regions but intervenes when unsafe trajectories are predicted. Experiments showed how different control modules influence patient motion, laying the groundwork for user studies on optimal interaction strategies.

Chapter 6 augments objective biomechanics with subjective patient feedback by introducing discomfort maps. Using a simple push-button interface, patients continuously reported discomfort during robotic therapy, allowing automatic reconstruction of their perceived safe and unsafe zones. This proof-of-concept demonstrates the feasibility of integrating subjective perceptions with biomechanical safety metrics for more patient-centered care.

Finally, the discussion reflects on the advancements proposed and their integration, while acknowledging limitations and outlining directions for future research. The thesis contributes a foundation for rehabilitation robotics that is sensitive not only to movement, but to the human body itself—an essential step toward safer, smarter, and more widely adopted rehabilitation technologies. Combining objective biomechanical metrics with subjective measures such as comfort could lead to therapy that is not only effective but also truly patient-centered.

Beyond the scientific and technical results, this work directly targets compelling societal challenges. Populations worldwide are aging, and demand for rehabilitation is growing faster than the supply of trained physiotherapists. Technologies that can enhance therapy safety, efficiency, and personalization are urgently needed to improve quality of life for patients, shorten recovery times, and reduce costs. At the same time, the integration of computational biomechanics into robotics opens avenues beyond healthcare: for instance, ergonomic workplace design, athletic training, or technologies that can assist us at home or in industrial settings.

Samenvatting

Musculoskeletale aandoeningen behoren wereldwijd tot de belangrijkste oorzaken van pijn, invaliditeit en verlies van onafhankelijkheid. Ze treffen miljoenen mensen, en de incidentie neemt aanzienlijk toe met de leeftijd.

Een van de meest voorkomende musculoskeletale blessures is een gescheurde rotator cuff van de schouder. Omdat deze musculotendineuze structuren anatomisch beperkt zijn tot een zeer nauwe ruimte tussen de schouderbladen, zijn ze vaak onderhevig aan trauma of slijtage. De behandeling van deze blessures is dringend, zowel medisch als sociaal: ze belemmeren de dagelijkse activiteiten, beperken de mogelijkheid om te werken en te sporten, en leiden tot hoge persoonlijke en zorgkosten. Revalidatie is essentieel voor herstel, maar is vaak langdurig en veeleisend voor zowel fysiotherapeuten als patiënten. Bovendien is de kans op recidief groot, omdat fysiotherapeuten niet over kwantitatieve instrumenten beschikken om de ontwikkeling van de complexe musculoskeletale structuren tijdens de therapie te monitoren.

In deze context biedt het gebruik van robotica mogelijkheden om handmatige patiëntmanipulatie te ondersteunen en geavanceerde sensoren te gebruiken om de patiënt te monitoren. Ondanks vooruitgang in robotontwerp en -besturing negeren huidige systemen echter nog steeds de biomechanica van de patiënt en zijn ze niet in staat om schadelijke belasting van regenererend weefsel te monitoren of te voorkomen. Dit proefschrift vult deze cruciale kennis gebrek door geavanceerde musculoskeletale modellen te integreren in de besturing van revalidatierobots. Door nieuwe algoritmen te ontwikkelen, is realtime schatting van diepe spieractiviteit en peesvervorming in de schouder mogelijk tijdens fysieke mens-robotinteractie. Door verbeterde biomechanische simulaties uit te breiden naar de integratie ervan in robottherapie, brengt dit werk de stand van zaken aanzienlijk en biedt het een coherent raamwerk voor biomechanisch geïnformeerde robotfysiotherapie.

Ten eerste bespreekt hoofdstuk 1 het werk dat in het huidige wetenschappelijke en klinische landschap is verricht, waarbij de beperkingen van bestaande revalidatiebenaderingen worden benadrukt en de integratie van biomechanica in robotica wordt gerechtvaardigd.

Hoofdstuk 2 introduceert de Rapid Muscle Redundancy (RMR)-solver, een nieuw algoritme voor het snel schatten van fysiologisch consistente spieractivaties in de schouder met behulp van een computationeel model. Door de glenohumerale stabiliteit te verbeteren en passieve vezelbijdragen te integreren, toonde de solver aan dat de activiteit van de rotator cuff niet alleen kan worden afgeleid uit oppervlakkige spieren, wat het belang van simulaties voor het monitoren van diepe spieractivatie onderstreept.

Hoofdstuk 3 bouwt voort op deze aanpak met BATON (Biomechanics-Aware Trajectory Optimization for Navigation of tendon strain maps), dat de RMR-solver integreert in robotbesturing. BATON maakt online schatting van peesbelasting mogelijk tijdens

mens-robotinteractie en genereert veilige en adaptieve revalidatiepaden die het risico op recidief verminderen.

Hoofdstuk 4 gaat verder in op onze aanpak voor het schatten van peesbelasting-skaarten, die vastleggen hoe kinematica, externe krachten en spieractivaties gezamenlijk de peesbelasting beïnvloeden. We laten zien hoe deze kaarten snel het effect van de weerstand van de robot op menselijke beweging kunnen beoordelen door de invloed van oefeningen te visualiseren, wat ze tot een nieuw diagnostisch en personalisatie-instrument voor fysiotherapeuten maakt.

Hoofdstuk 5 richt zich op revalidatiefasen waarin patiënten worden aangemoedigd om actief te bewegen, maar kwetsbaar blijven voor hernieuwde blessures. Het introduceert een gedeelde besturingsstrategie waarbij de robot vrije beweging in veilige zones toestaat, maar ingrijpt wanneer gevaarlijke trajecten worden voorspeld. Onze experimenten hebben aangetoond hoe verschillende besturingsmodules de beweging van de patiënt beïnvloeden, wat de basis legt voor gebruikersonderzoek naar optimale interactiestrategieën.

Hoofdstuk 6 combineert de objectiviteit van biomechanica met subjectieve feedback van de patiënt door de introductie van ongemakskaarten. Met behulp van een eenvoudige knopinterface kunnen patiënten continu ongemak rapporteren tijdens robottherapie, wat de automatische reconstructie van waargenomen veilige en gevaarlijke zones mogelijk maakt. Dit proof of concept toont aan dat het mogelijk is om subjectieve percepties te integreren met biomechanische veiligheidsmaatregelen voor meer patiëntgerichte zorg.

Tot slot reflecteert de discussie op de voorgestelde ontwikkelingen en hun integratie, waarbij de beperkingen ervan worden erkend en richtingen voor toekomstig onderzoek worden geschetst. Het proefschrift legt de basis voor revalidatierobotica die niet alleen gevoelig is voor beweging, maar ook voor het menselijk lichaam zelf: een fundamentele stap naar veiligere, slimmere en potentieel geavanceerdere revalidatietechnologieën. De combinatie van objectieve biomechanische metingen met subjectieve factoren zoals comfort kan leiden tot therapie die niet alleen effectief is, maar ook echt patiëntgericht.

Naast wetenschappelijke en technische prestaties richt dit werk zich direct op urgente maatschappelijke uitdagingen. De wereldbevolking vergrijst en de vraag naar revalidatie groeit sneller dan het aanbod van gekwalificeerde fysiotherapeuten. Technologieën die de veiligheid, efficiëntie en personalisatie van therapie kunnen verbeteren, zijn dringend nodig om de kwaliteit van leven van patiënten te verbeteren, de hersteltijd te verkorten en de kosten te verlagen. Tegelijkertijd opent de integratie van computationele biomechanica in robotica mogelijkheden buiten de gezondheidszorg: bijvoorbeeld in ergonomisch werkplekontwerp, sporttraining of technologieën die ons thuis of in de industrie kunnen ondersteunen.

Sommario

Le lesioni muscoloscheletriche sono tra le principali cause di dolore, disabilità e perdita di indipendenza nella popolazione mondiale. Colpiscono milioni di persone, con un'incidenza che aumenta rapidamente con l'età. Una delle lesioni muscoloscheletriche più comuni è la rottura, parziale o completa, dei tendini che costituiscono la cuffia dei rotatori della spalla. Poiché questi tessuti muscolo-tendinei sono anatomicamente confinati in uno spazio molto ristretto tra le ossa della spalla, sono frequentemente soggetti a traumi o usura. Il trattamento appropriato di queste lesioni è un tema urgente sia dal punto di vista medico che sociale: esse infatti compromettono le attività quotidiane, limitano la capacità di lavorare e praticare sport, e generano alti costi sociali e sanitari. Una corretta riabilitazione è essenziale per il recupero, ma spesso è lunga e richiede un impegno consistente sia da parte dei fisioterapisti sia dei pazienti. Inoltre, è soggetta a ricadute come il reinfortunio, poiché i fisioterapisti non dispongono di strumenti quantitativi per monitorare come l'apparato muscoloscheletrico del paziente risponde alla terapia.

In questo contesto, l'adozione di dispositivi robotici offre importanti opportunità per supportare la mobilizzazione del paziente e fornire sensori sofisticati per monitorarlo. Tuttavia, nonostante i progressi nel design e nel controllo dei robot, i sistemi attuali non sono in grado di considerare l'apparato muscoloscheletrico del paziente e, pertanto, non possono monitorare o prevenire carichi dannosi sui tessuti interessati. Questa tesi affronta tale lacuna di conoscenza incorporando modelli muscoloscheletrici all'avanguardia all'interno degli algoritmi di controllo dei robot riabilitativi. Attraverso lo sviluppo di nuovi algoritmi, il nostro lavoro permette la stima in tempo reale dell'attività dei muscoli profondi e del carico meccanico sui tendini della spalla durante l'interazione fisica uomo-robot. Il contributo di questa tesi spazia da simulazioni biomeccaniche avanzate fino alla loro integrazione nell'esecuzione della terapia robotica, avanzando significativamente lo stato dell'arte, e creando un quadro coerente per la fisioterapia robotica basata sulla biomeccanica.

In primo luogo, il Capitolo 1 contestualizza il nostro lavoro nel panorama scientifico e clinico attuale, evidenziando i limiti degli approcci riabilitativi esistenti e motivando l'integrazione della biomeccanica nella robotica.

Il Capitolo 2 introduce il Rapid Muscle Redundancy (RMR) solver, un nuovo algoritmo che utilizza un modello computazionale per stimare rapidamente il livello di attivazione dei muscoli della spalla. Vincolando la stabilità gleno-omerale e considerando il contributo passive delle fibre muscolotendinee, il solver dimostra che l'attività della cuffia dei rotatori non può essere dedotta solo dai muscoli superficiali, sottolineando quindi la necessità di simulazioni per monitorare l'attivazione dei muscoli profondi.

Il Capitolo 3 sviluppa BATON (Biomechanics-Aware Trajectory Optimization for Navigation of tendon strain maps), che integra il nostro algoritmo RMR nel controllo robotico. BATON consente di stimare l'elongazione dei tendini durante l'interazione uomo-

robot in tempo reale, generando traiettorie riabilitative sicure e adattive che riducono il rischio di reinfortunio.

Il Capitolo 4 approfondisce ulteriormente il nostro approccio per lo sviluppo di “mappe del carico tendineo”, che rappresentano come cinematica, forze esterne e attivazioni muscolari influenzino congiuntamente tale variabile. Questo capitolo illustra come queste mappe possano stimare rapidamente l’effetto della resistenza robotica al movimento umano visualizzando l’effetto di vari esercizi sul carico tendineo della spalla. Tali mappe possono anche essere impiegate dai fisioterapisti, offrendo loro un nuovo strumento diagnostico e di personalizzazione della terapia.

Il Capitolo 5 affronta la fase riabilitativa in cui i pazienti sono incoraggiati a muoversi attivamente ma rimangono vulnerabili al reinfortunio. Introduce strategie di controllo condiviso in cui il robot permette al paziente di effettuare movimenti liberi nelle aree sicure ma interviene quando vengono rilevate traiettorie pericolose. I nostri esperimenti dimostrano come diverse variazioni del livello di correzione robotica influenzino il movimento del paziente, gettando le basi per trovare le migliori strategie di interazione con tra robot e pazienti.

Il Capitolo 6 integra le metriche oggettive proprie della biomeccanica con il feedback soggettivo del paziente introducendo “mappe di scomodità”. Utilizzando una semplice interfaccia equipaggiata con un pulsante, i pazienti possono riportare continuamente il loro livello di scomodità o dolore durante la terapia robotica, permettendo la ricostruzione automatica delle zone da loro percepite come sicure o dolorose. I nostri esperimenti preliminari dimostrano la fattibilità dell’integrare percezioni soggettive con metriche di sicurezza biomeccanica per una cura più attenta alle varie dimensioni del paziente.

Infine, proponiamo al lettore una discussione che riflette sui progressi proposti e sulla loro integrazione, riconoscendo i limiti e delineando direzioni per la ricerca futura. In conclusione, il lavoro presentato in questa tesi contribuisce a creare una base per la robotica riabilitativa sensibile non solo al movimento, ma al corpo umano stesso—un aspetto essenziale per tecnologie riabilitative più sicure, intelligenti e accessibili. Combinare metriche biomeccaniche oggettive con misure soggettive come la comodità dei movimenti fisioterapici potrebbe portare a terapie non solo efficaci ma veramente centrate sul paziente.

Oltre ai risultati scientifici e tecnici, questo lavoro affronta direttamente sfide sociali rilevanti. La popolazione occidentale sta invecchiando e la domanda di riabilitazione cresce più rapidamente della disponibilità di fisioterapisti qualificati. Tecnologie in grado di migliorare la sicurezza, l’efficienza e la personalizzazione della terapia sono urgentemente necessarie per migliorare la qualità della vita dei pazienti, ridurre i tempi di recupero e abbattere i costi. Allo stesso tempo, l’integrazione della biomeccanica computazionale nella robotica apre possibilità oltre l’ambito sanitario: ad esempio nella progettazione ergonomica dei luoghi di lavoro, nell’allenamento sportivo o in tecnologie che possano assisterci sia casa che sul luogo di lavoro.

1

Introduction

This initial chapter positions this dissertation in the current broader societal context, motivating the need for new tools to augment traditional physiotherapy. It identifies the opportunities offered by combining precise motion control and sensing offered by robots with insights that can be gained from biomechanical simulations, and presents the specific research gaps addressed in later chapters. Shoulder rotator cuff tendon tears are introduced as prominent musculoskeletal injuries that could benefit from new treatment techniques that only a combination of robotics and biomechanical simulation can unlock.

1.1. Motivation

Progressive population aging, persistence of strenuous manual labor, and widespread engagement in sports are leading to an increase of musculoskeletal injuries in Western societies (Fig. 1.1) [1–3]. These injuries, such as back pain, osteoarthritis, and shoulder rotator cuff tears, affect the human muscular or skeletal system, leading to decreased independence and quality of life, as well as reduced productivity at work and earlier retirement [4]. It is estimated that around 30% of Europeans were living with the consequences of a musculoskeletal injury in 2019 [5], with this number rising to 1.7 billion people worldwide [6]. From an economic point of view, these injuries accounted for combined healthcare expenditures and productivity losses of about €240 billion in Europe [7], and beyond $940 billion in the USA [1].

After a musculoskeletal injury, rehabilitation programs play an important role in restoring body functions. In particular, early interventions and therapy volumes are often correlated with shorter recovery times [8, 9], highlighting the important role of physiotherapists (PTs) for successful rehabilitation. However, shortages among PTs, combined with increasing rehabilitation demands, are intensifying the strain on our healthcare systems, and this will continue into the foreseeable future [10, 11]. Fewer physiotherapists will have less time to handle each of their patients, leading to lower therapy intensity and longer recovery time.

Beyond the growing need for rehabilitation, maintaining high therapy intensity can be physically demanding for PTs themselves [12]. Prolonged manual manipulation of their patients is standard practice in the early stages of rehabilitation, to provide external guidance and movement support (Fig. 1.2, right) [13, 14]. Overall, physical fatigue and time limit the number of patients that can be handled by the same PT.

Another issue that complicates the treatment of musculoskeletal injuries is the intrinsic complexity of the human body. Since PTs generally cannot measure “what happens under the patient’s skin”, they are unable to monitor the effect of their therapy in a quantitative manner.

Figure 1.1: *Aging population (A) [15], demanding manual work (B) [16], and extreme engagement in sport (C) [17] are leading causes of musculoskeletal injuries nowadays.*

Figure 1.2: *Robots have already been introduced in clinics for aiding neuro-rehabilitation tasks, such as gait training (left) [18]. In many cases, however, prolonged manual manipulation of the patient amounts to a significant part of physiotherapists' daily activities (right) [19]. Robots that are better suited to assist in musculoskeletal rehabilitation could be new tools in the hands of therapists, reducing their fatigue and providing new patient monitoring opportunities.*

1.2. Potential of Robotics

The challenges outlined above—limited therapy dosage, therapist fatigue, and lack of monitoring tools—highlight the urgent need for technological solutions to support musculoskeletal rehabilitation. In this context, robotic devices can play a key role in enhancing traditional physiotherapy. Robots can interact physically with patients to guide and support their movements, simultaneously increasing therapy dosage and assisting PTs to relieve them from excessive fatigue. In addition, robotic devices may also be equipped with sensors to better monitor patients' poses and forces, and they may accurately control the training of specific motions.

Thanks to these benefits, robots have already been employed to support the neurorehabilitation of stroke survivors, gaining popularity in both clinics and academia [20]. A prominent example is gait re-training (Fig. 1.2, left) [21], where numerous robotic exoskeletons have been developed. More recently, systems targeting the full mobility of the upper limbs have also become available, both in research and commercially (e.g., ANYexo [22], Harmony [23], and ArmeoPower [24]). These robots and their control algorithms typically provide functional training, often simulating activities of daily living to help patients regain independence. To stimulate neuroplasticity and promote motor learning [25, 26], they emphasize repeated movements, enriched through haptic feedback [27] and virtual reality [28, 29]. Building on these development, robots specifically dedicated to automate repetitive movements in physiotherapy have been recently introduced on the market (e.g., ROBERT [30]). However, approaches designed for neurological patients may fall short when transferred to patients recovering from musculoskeletal injuries.

The sensory information gathered by robots could be used to further enhance physiotherapy treatments. In particular, data about physical human-robot interaction

(pHRI) may unlock monitoring of patients' musculoskeletal metrics. Once available, these metrics could be incorporated into the robots' control algorithms, to precisely regulate patient movement while reducing PT fatigue. Yet, this potential to fundamentally improve both robotic and human understanding of the patients' musculoskeletal system has remained largely unexplored until now.

1.2.1. Enabling biomechanics-aware robots

Robots that are aware of how their physical interaction with a patient affects the underlying human biomechanics could be invaluable tools to enable PTs to treat patients safely and with better outcomes. A promising option to retrieve the necessary musculoskeletal information comes from computational biomechanics, which, in the past two decades, has proven capable of explaining why humans move in the way they do, and even predicting their pathological behavior (Fig. 1.3) [31–34]. Moreover, biomechanical simulations allow one to monitor metrics that are not easily accessible (or not accessible at all) via sensors, such as surface electromyography (EMG). Examples of such metrics are joint loading, tendon strain, and activations of muscles lying deep under the human skin, all of which are currently unavailable to therapists, but are potential measures to evaluate and modulate robotic assistance to improve therapy outcomes and avoid re-injury [35]. As such, augmenting robotic physiotherapy with high-fidelity digital simulations of the human body holds great potential to improve musculoskeletal treatment. A similar lesson comes from many fields in engineering, such as aviation: early airplanes relied mostly on intuition, but true safety and performance came only with simulation-based insights.

In fact, the robotics community is still new to the use of biomechanical metrics from high-fidelity musculoskeletal models. This is due to the complexity of integrating musculoskeletal models and their metrics into the control of a robotic system, and to their traditionally heavy computational cost. When developing control strategies for rehabilitative robots, research has mostly focused on other metrics, like contact forces [36], workspace reachability [37], or similarity to human therapy [38]. Simplified torque-driven human models have also been used to monitor human kinematics and generalized joint torques [39], while the few applications of more realistic models have been limited to offline analysis [40, 41]. However, only high-fidelity models of the human musculoskeletal system can deliver deeper insights into the inner functioning of the human body and improve robotic-assisted physiotherapy [35]. The time appears ripe to take the next step towards the integration of musculoskeletal metrics into robotic control, to enable adaptation of pHRI to the patient response.

Moving in this direction, very recent work from Sartori et al. leveraged EMG-tracking musculoskeletal simulations to control an assistive back exoskeleton in real-time [45–47]. Despite these first promising results, online use of high-fidelity biomechanics has not yet been employed to control robots for physical therapy. Such an application has great potential, as it can enable robots to deliver precisely the kind of rehabilitation that each patient needs, adjusting therapy in response to the patient's actions. For example, a robotic device could control the evolution of tissue and joint loading after a muscle-tendon tear to keep loads within healthy bounds, or modulate rehabilitation

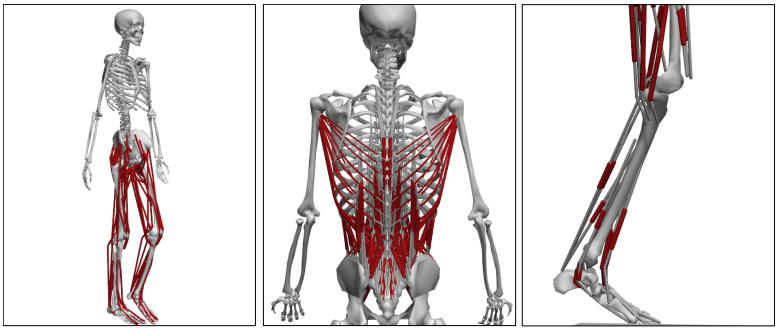

Figure 1.3: *Musculoskeletal models have been developed to advance our understanding of how we move. Human movement can be simulated with full-body models (left) [42], or by focusing on specific body parts for more detailed analysis: models of the human spine may uncover hidden reasons for back pain (center) [43], while analysis of the internal forces generated at the knee might unveil treatments for osteoarthritis (left) [44].*

exercises to target specific muscle groups.
Stemming from the observations above, the work presented in this thesis advances the integration of biomechanical modeling into closed-loop robot controllers to increase the safety and effectiveness of robotic-assisted physical therapy. Specifically, these improvements are presented with regard to shoulder rehabilitation, as detailed in the following paragraph.

1.3. Focus: rotator cuff tears

The human shoulder is a very complex biological mechanism that permits the largest range of motion in the human body. During movement and at rest, the shoulder is stabilized by the rotator cuff muscles that span the glenohumeral joint (between the humerus and the clavicle, see Fig. 1.4) and hold the arm in place. Unfortunately, tendons in the rotator cuff are susceptible to injuries and tears due to repeated overloading or traumatic events. Studies have found the prevalence of rotator cuff tears to be more than 22% in the general population, and as high as 50% in people over 65 years of age [48, 49]. Injuries to the rotator cuff tendons are one of the primary causes of pain and disability in the upper limbs, and may result in the inability to perform many daily living tasks, leisure activities, and occupations [50–52].
Rotator cuff tears often require lengthy rehabilitation, sometimes preceded by surgical interventions, to restore shoulder functionality. The early stages of rotator cuff rehabilitation are particularly challenging and labour-intensive for the therapists, as they typically involve passive mobilization where the therapist manually guides the patient's shoulder through specific movements [53, 54]. These early interventions aim to restore the range of motion and prevent stiffening of the joint structures. However, they also carry a high risk of tendon re-injury, particularly when healing tissues are exposed to excessive loading [55].
Re-injuries during early rehabilitation are often linked to unsafe arm movements or

unintended muscle activation by the patient, leading to increased mechanical loading on the tendons.
In this context, tendon strain emerges as an important metric to consider for minimizing re-injury [56]. Yet, direct measurement of tendon strain during therapy is infeasible, since tendons lie deep beneath other tissues and bones. As a consequence, only biomechanical models could reveal the musculoskeletal response to robot-assisted movement, to unlock continuous and non-invasive monitoring of tendon strain both during passive and more active exercises (Fig. 1.5). In this thesis, I aim to harness state-of-the-art biomechanical simulations to improve robotic physiotherapy and develop an approach to robotic-assisted rotator cuff rehabilitation that can enhance safety and efficacy by actively shaping rotator cuff tendon strain.

1.4. Aims

The overarching aim driving the research in this thesis is to design and develop a framework for robotic-assisted physiotherapy that can leverage quantitative metrics of the patient's musculoskeletal system to improve rehabilitation of the rotator cuff tendons. This central aim can be summarized by the following research question:

"How can musculoskeletal model–derived metrics be integrated into rehabilitation robot control online, shaping robot-mediated movement to prevent potentially harmful loading of rotator cuff tendons?"

To explore this broad question, the thesis is structured around five specific sub-questions, each treated in a dedicated chapter. These sub-questions address existing gaps in the state of the art, reflecting the different challenges for enabling physiotherapy robots to interpret, act upon, and adapt to biomechanical information throughout various stages of rehabilitation. The solutions that will be proposed are applicable beyond robotic rehabilitation: for instance, while rapid estimation of muscle activations in this thesis serves to infer tendon strain levels for rotator cuff rehabilitation, the same capability can enable broader human movement analysis.

Q1: *What musculoskeletal modeling and computational strategies can enable rapid and accurate estimations of human muscle activations during shoulder movement?*

Q2: *How can a robotic system fuse data from physical human-robot interaction and musculoskeletal simulations to estimate and navigate rotator-cuff strain online during early-stage shoulder rehabilitation?*

Q3: *How do model-based strain estimates for various rotator cuff tendons change as a consequence of robot-administered shoulder physiotherapy?*

Q4: *In what ways can a physiotherapy robot limit active human movements within safe tendon strain bounds leveraging a shared control approach?*

Q5: *How can direct human discomfort perception be captured to further personalize robotic-assisted shoulder rehabilitation treatment?*

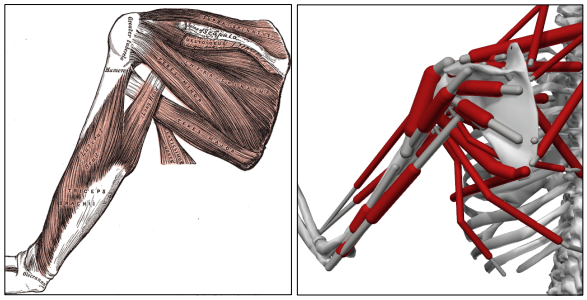

Figure 1.4: *The rotator cuff muscles stabilize the shoulder and contribute to its mobility. However, their complex anatomy (left) [57] complicates understanding of internal tissue function and loading, making the use of musculoskeletal models essential (right) [58].*

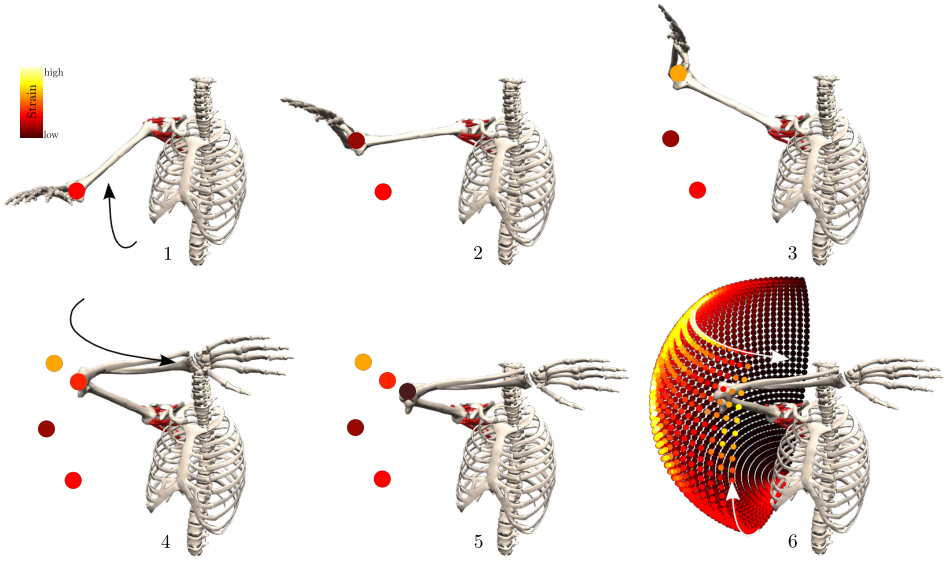

Figure 1.5: *A musculoskeletal model carries information about how physical quantities evolve during human movement. For example, a shoulder model can be posed in various configurations to extract maximum strain across the rotator cuff tendons, producing state-dependent maps of the strain.*

1.5. Contributions and Thesis Outline

My main contributions to the field are presented in the next two chapters. Jointly, they enable (1) rapid estimation of human muscle activations through a musculoskeletal model, leveraging novel formulations of the underlying physiological constraints [59], and (2) embed the resulting online model-based strain estimates into a biomechanics-aware controller to shape robot-mediated therapy [60]. Following the order of the sub-question above, subsequent chapters extend these contributions with (3) more extensive

analyses of tendon strain during various robotic-assisted exercises, (4) strategies for safe active rehabilitation, and (5) inclusion of patient comfort into therapy design. Below, I summarize briefly each chapter in relation to the research (sub)question it answers, while a synoptic view of the outline of my thesis can be found in Fig. 1.6.

- Chapter 2 improves current computational methods (Q1) by presenting a novel algorithm to estimate physiological muscle activations in a musculoskeletal model: the **R**apid **M**uscle **R**edundancy (RMR) solver. Compared to previous state-of-the-art solutions, this solver achieves improved performance by offering more physiologically plausible solutions at a reduced computational cost. Such characteristics make the RMR solver an indispensable building block towards rapid and accurate estimation of biomechanical metrics for robotic control.

- Chapter 3 addresses the integration of online biomechanics into robotic control (Q2), which is the core of this thesis. In this chapter, I introduce BATON, a novel approach to **B**iomechanics-**A**ware **T**rajectory **O**ptimization for robotic **N**avigation of rotator-cuff strain. BATON embeds a musculoskeletal model of the shoulder into a receding horizon controller and uses the RMR solver to reveal how pHRI affects internal tissue loading. With this information, BATON can replan rehabilitation movements in real-time, adapting to unpredictable voluntary or reflexive human actions during therapy.

- Chapter 4 further supports the need to consider muscle activations when estimating rotator cuff tendon strain during robotic-assisted movement (Q3). Four conventional therapy exercises are performed by a subject with the assistance of our robotic system. We post-processed the robot data to elucidate variations in tendon strain resulting from increased muscle activations. Significant changes are observed, strengthening the need for systems like BATON, capable of modifying therapeutic interventions in real-time to protect patients from unsafe scenarios.

- Chapter 5 shifts the focus to guaranteeing safety also in later stages of rotator cuff therapy, when patients gradually gain autonomy and are expected to take a more active role (Q4). In this sense, we developed two shared control algorithms, each aiming to deflect unsafe human movements towards low-strain shoulder poses. These two algorithms are compared during experiments in our lab with a healthy participant, paving the way to wider user studies in the future.

- Chapter 6 augments computational biomechanics, introducing the concept of patient-specific “pain-maps”, to accurately capture human discomfort during pHRI (Q5). As patient comfort is expected to be fundamental for user acceptability, but cannot yet be effectively predicted, we propose a system to log personalized discomfort information. I discuss how this information can be fused into the other methods presented in this thesis to augment them toward a holistic approach to personalized rehabilitation.

- Chapter 7 discusses the advancements developed in this thesis, and draws directions for further research.

The code and data associated with this thesis are made available as a dataset on the *4TU.ResearchData* platform[1].

1.6. Studies Enabled by this Thesis

Next to the work in this thesis, other collaborative research projects were enabled and pursued jointly. The following gives a brief overview of these research directions.

Patient- and Teleoperator-led Robotic Physiotherapy: in a collaborative work with Stephan Balvert [61], we implemented a first algorithm to protect patients from excessive tendon strain during shoulder rehabilitation. Given a strain threshold, the algorithm automatically identifies areas across the human range of motion where safe movements are possible, and establishes virtual barriers around unsafe zones through our robotic system. In this way, patient movement is limited when the given tendon threshold is exceeded, both during autonomous patient movement and during tele-rehabilitation. Chapter 5 in this thesis expands this initial approach to anticipate unsafe movements.

Model-based analysis of bench press exercises: in collaboration with Lisa Noteboom (VU Amsterdam), we used the Rapid Muscle Redundancy solver (presented in Chapter 2) to analyze the dependency between various bench press techniques and shoulder loading [62]. We analyzed data from 10 experienced strength athletes and sought differences in glenohumeral contact forces and rotator cuff muscles activation resulting from changes in grip width, scapula pose, and medio-lateral exerted barbell forces. Overall, we found that scapula retraction and grip widths had a significant correlation with decreased glenohumeral posterior shear force and rotator cuff activity. These findings, achieved through our musculoskeletal simulation pipeline and in line with expected results from physiological literature, underscore the utility of these computational techniques also for rehabilitation robotics.

Experimental validation of predicted glenohumeral contact forces: with Ibrahim Hasan (KTH), we compared the predictions of the Rapid Muscle Redundancy solver to data collected from a participant with an instrumented shoulder implant [63]. In particular, we validated the magnitude and direction of the predicted reaction force at the glenohumeral joint against the experimental data. We found good agreement for various mathematical implementations of the glenohumeral joint stability constraint in our algorithm, further demonstrating its broad applicability.

[1] https://doi.org/10.4121/011863e8-86d0-4894-91f5-d580c5cade27

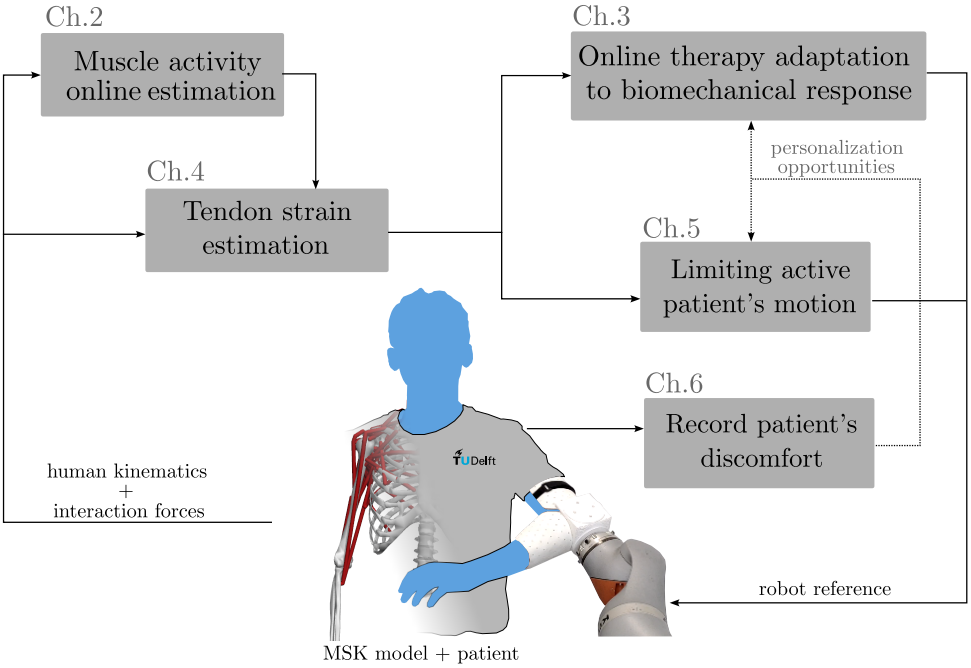

Figure 1.6: *Outline of this thesis. The novel muscle estimation algorithm proposed (RMR solver) is used to inform the online trajectory adaptation of early robotic-assisted rotator cuff therapy (BATON), and to thoroughly inspect the changes that physical human-robot interaction induces on tendon strain. A shared control solution is proposed to guarantee safety during active patient movements as therapy progresses, and capturing patient discomfort could improve and personalize robotic therapy as a whole.*

2

Estimating shoulder muscle recruitment via a rapid solver

Shoulder movements in activities of daily living and rehabilitation exercises are produced by the coordinated action of large superficial muscles and deeper stabilizers spanning multiple degrees of freedom. Tissue loading is strongly influenced by muscle recruiting, making accurate monitoring of muscle activations essential for safe rehabilitation after musculoskeletal injury. However, activations of deep muscles such as the rotator cuff cannot be directly measured, limiting our understanding of their role and raising the risk of re-injury during post-tear physiotherapy.

Musculoskeletal simulations offer a powerful way to infer the function of these deep muscles. Yet, existing approaches often neglect key physiological requirements—most notably, stability of the glenohumeral joint (between humeral head and scapula). Moreover, it remains unclear whether enforcing these constraints meaningfully changes predicted activations, as no systematic evaluation had been conducted.

We address this gap by examining how glenohumeral stability influences estimated shoulder muscle activity. To this end, we introduced the Rapid Muscle Redundancy (RMR) solver to efficiently include constraints on joint reaction forces (JRFs) from a musculoskeletal model while accounting for activation dynamics and passive fiber forces. Its efficiency allowed us to test over 3600 motion trials within experimental uncertainty, comparing results with and without GH stability enforced. Our analysis showed that enforcing stability significantly increased estimated rotator cuff activity while leaving most superficial muscles unaffected, making EMG validation from superficial muscles alone insufficient to infer rotator cuff activity.

By enabling physiologically consistent and rapid estimation of muscle recruitment, the RMR solver opens new opportunities for studying upper-limb function and for applications where fast and reliable tissue load estimation is crucial (e.g. real-time decision-making in robotic-assisted shoulder rehabilitation).

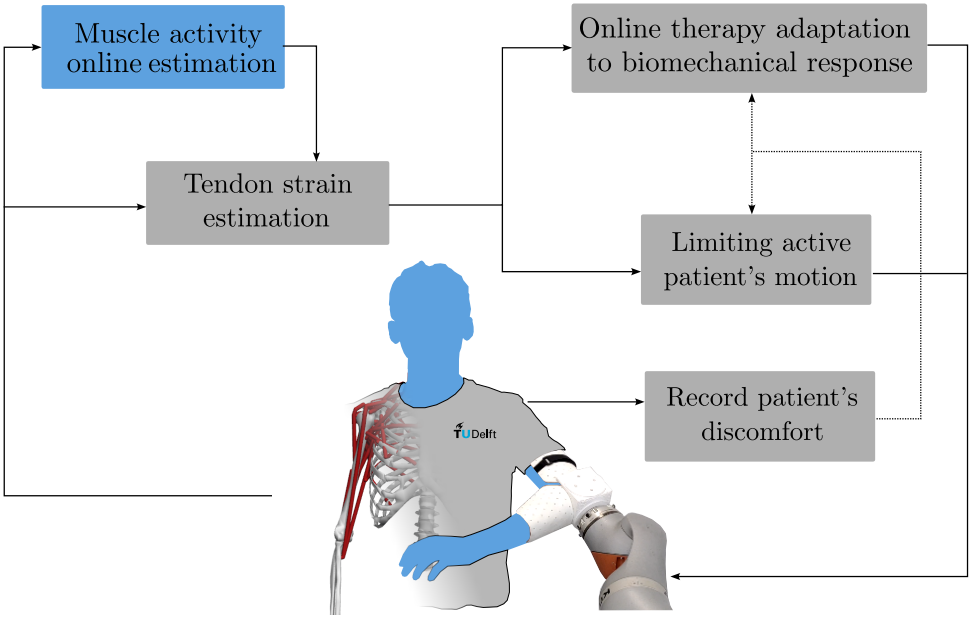

This chapter is based on **I. Belli**, Sagar Joshi, J. M. Prendergast, I. Beck, C. Della Santina, L. Peternel, and A. Seth. Does enforcing glenohumeral joint stability matter? A new rapid muscle redundancy solver highlights the importance of non-superficial shoulder muscles. PLOS One 18.11 (2023): e0295003. [59].

A preliminary version of this work, "A rapid muscle redundancy solver for estimating shoulder loading", was finalist for the Andrzej Komor New Investigator Award at the XIX International Symposium on Computer Simulation in Biomechanics July 26th–28th 2023, Kyoto.

2.1. Introduction

Health and proper mobility of the shoulder are important in performing daily activities and maintaining independence, as the shoulder enables one to lift objects, groom and get dressed, or play a sport [64]. The mobility of the shoulder permits the largest range of motion among human joints, thanks to a mechanism composed of many bone segments, joints, and muscles. This impressive mobility comes at the cost of reduced stability of the glenohumeral joint (GH), which does not rely on any bony "socket", but is instead stabilized by soft tissues including muscles that span the joint [65, 66]. To preserve glenohumeral joint integrity, a balance must be reached by soft tissue forces and the contact force generated at the glenoid [67]. A subset of shoulder muscles, collectively known as the rotator cuff muscles (infraspinatus, supraspinatus, subscapularis and teres minor), are considered glenohumeral stabilizing muscles [68], which are of particular interest, as they are a common source of shoulder injuries [48]. However, accessing directly rotator cuff muscles is particularly difficult via surface electromyography (EMG), as they lie deeper below other shoulder muscles. In this context, the use of musculoskeletal models of the shoulder can be a powerful tool to gain insights into rotator cuff muscle function.

For musculoskeletal models that represent the complete musculature, it is challenging to predict the muscle forces required to generate specific motions, because human (and animal) joints typically feature more actuators (i.e., muscles) than movement degrees of freedom (DoFs). As a consequence, there are infinite possible combinations of muscle forces that can generate the same joint torque and acceleration, as many muscles are mechanically redundant, leading to the so-called muscle redundancy problem. This is also true for the human shoulder, which is actuated by a combination of larger surface and deeper rotator cuff muscles that span multiple joints of the shoulder complex.

Several optimization-based approaches have been employed to solve the muscle redundancy problem in biomechanical simulations, which fall into three categories: methods integrating the dynamics of the model [69–73], methods considering the model statically [74–78], and data-driven approaches [79, 80]. However, despite these efforts, there are several open issues regarding the estimation of muscle forces from a musculoskeletal model [81]. In particular, methods that integrate the dynamics of the model (such as the popularly employed CMC [70]) permit respecting physiological constraints on the activation dynamics, but are computationally expensive. Bypassing the direct numerical integration of the system dynamics with solutions like direct collocation lightens this burden but remains prohibitively slow for real-time applications [73].

On the contrary, static methods disregarding dynamic constraints and considering each instant of the movement as independent are fast but could lead to non-physical solutions, so bounds are employed to limit the evolution of the optimization variables between consecutive time-steps [76, 77, 82]. Nonetheless, the widely used OpenSim's implementation of the static optimization approach [78] disregards the activation dynamics, and does not account for the effects of passive fiber forces. Neglecting passive forces leads to simplifications of muscle function [83], together with poor performances in estimating antagonist muscle activity at the GH joint [84]. Finally, recent data-driven machine learning methods achieved promising results [79, 80], yet

they currently disregard musculoskeletal properties altogether, retaining little direct connection with the way the human body actually functions.
Altogether, the methods described above do not allow the inclusion of constraints on the joint reaction forces (JRFs) arising during the movements. A recently-developed open-access framework could enhance them to account for stability of the joint [85], but unfortunately its formulation still ignores the passive contribution of the muscle fiber when estimating the JRFs, and computations to guarantee stability remain quite costly. In view of these reasons, a number of previous upper-extremity simulation studies concerned with shoulder function have disregarded the GH stability issue [58, 86–88]. Other studies, given the biomechanical importance of joint stability, have overcome these limitations and included constraints on the JRF at the glenoid when estimating individual muscle forces in the shoulder [72, 89–96], or during forward dynamics simulations of shoulder movements [97]. However, most of these works have not compared how GH stability affects the estimates of individual shoulder muscle activity [72, 89, 91, 93, 94, 96], and recent investigations reported that constraining the GH force did not influence muscle activations [95].
Changes in estimated muscle activation when GH stability is enforced have been investigated during a box-lifting task [98], and significant differences were observed only in the supraspinatus. However, this analysis relied on OpenSim's Static Optimization algorithm [78], thus neglecting the passive contribution of the muscle fibers. Moreover, the formulation to enforce GH stability was not provided, and the implementation was not made available. A similar investigation was performed to understand the effects of GH stability on the estimates of muscles of the shoulder, but only in the case of serious simulated rotator cuff tears, and implementation was not released [90]. In summary, the effects of enforcing the GH stability on estimated muscle activity when resolving the muscle redundancy problem are not fully understood. Additionally, including constraints on the JRF at the glenoid has been done ad hoc and inefficiently, solving the multibody system for its reaction forces, and none of such methods is publicly available.
Therefore, the main aim of our study is to understand the effect of GH stability on rotator cuff muscles of the human shoulder from a musculoskeletal modeling approach. To do this, we developed an efficient musculoskeletal muscle activity/force solver to analyse human movement while respecting explicit physiological constraints on JRFs: we refer to this new method as the Rapid Muscle Redundancy (RMR) solver.

2.2. Methods

To understand the effects of GH stability on estimated muscle activity of the shoulder, we employed a musculoskeletal model of the shoulder [58] implemented in OpenSim [31, 99] with experimental data and developed the RMR solver to include constraints on the GH joint reaction force (JRF) in a computationally efficient manner. Estimated muscle activations were compared to EMG and estimates from CMC (Figure 2.1). Unlike other static optimization methods [74–78, 82], the RMR solver estimates muscle activations and joint forces via numerical optimization introducing three new features: first, muscle forces include active and passive forces that are length and velocity dependent; second, consistency of activation rates are enforced using linear constraints, and third, JRFs are

expressed as closed-form linear functions of the estimated muscle activation during the evaluation of the constraints. We compared our results against the state of the art and tested our conclusions for robustness against the uncertainty in motion data. We applied the RMR solver to estimate muscle activity during three shoulder tasks with and without GH stability to understand the effect of GH stability.

2.2.1. Rapid muscle redundancy (RMR) solver

To tackle the main aim and facilitate the study of the effect of GH stability on rotator cuff muscles of the human shoulder, we developed a novel Rapid Muscle Redundancy (RMR) solver. The objectives for the RMR development were:

- **O1**: efficiently solve the muscle redundancy problem in a musculoskeletal model,
- **O2**: constrain the solver to realistic joint reaction forces that will maintain GH stability,
- **O3**: respect physiological constraints that include passive fiber forces, and guarantee continuity of activation profile.

The RMR solver can estimate individual muscle forces and activations given an OpenSim musculoskeletal model and movement data. The solver solves the muscle redundancy problem efficiently (*O1*) by including constraints on joint reaction forces (*O2*), passive fiber force, and on the continuity of muscle activations (*O3*). The inputs to the RMR solver are the joint trajectories (angles, speeds, and their accelerations) from experimental data processing (detailed below), while the outputs are estimated muscle activations and JRFs at every corresponding time point of the joint trajectories (Figure 2.1). Like static optimization [78], the development of the muscle redundancy problem leverages numerical optimization at each time step of the input motions. The primary difference is that more physiologically realistic behavior (*O3*) can be enforced by means of including time-dependent constraints on reaction forces and on the rate of activation change, while also accounting for passive fiber forces.

At a given time instant t_k, a non-linear programming (NLP) problem is formalized and solved to retrieve the optimal combination of activations $\boldsymbol{a}_k \in \mathbb{R}^{N_\mathrm{m}}$ and controls $\boldsymbol{c}_k \in \mathbb{R}^N$ required to simulate the experimental motions, given the N_m muscles and N other actuators present in the model. We also consider the latter as, for example, it is common practice to include an ideal (reserve) actuator for each DoF ($N = N_\theta$) to enable the model to always achieve the experimental movement in the face of modeling and/or experimental data inaccuracies, where the reserves' magnitude is an error metric [100]. The model state, including joint coordinate (angle) values and speeds $\boldsymbol{\theta}_k$, $\dot{\boldsymbol{\theta}}_k \in \mathbb{R}^{N_\theta}$ for all the N_θ coordinates, is updated once per instant with kinematics determined from the IK analysis of the experimental motion. Given the state of the model, muscle-related variables such as length and lengthening speed derived from the model can be considered fixed (for that instant), which simplifies the formalization of the constraints to the NLP problem and enables the efficient solution of the muscle redundancy problem.

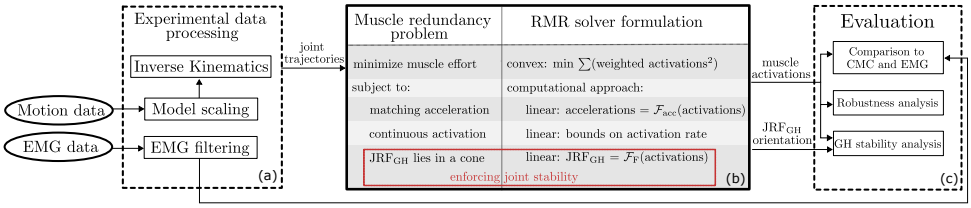

Figure 2.1: *Overview of the methodology to estimate and compare the effects of GH stability on shoulder muscle activity. An inverse analysis begins with the experimental motion and measured EMG (a), which is processed via an OpenSim model to determine joint trajectories from inverse kinematics. Joint trajectories are filtered and splined to estimate joint speeds and accelerations, which serve as inputs to the Rapid Muscle Redundancy solver (b). RMR exploits an optimization problem formulation with a convex minimum effort objective, and expresses acceleration, continuous activation, and JRFs as linear functions of the activations (as design variables). We tested estimated activations against CMC and EMG, as well as inherent uncertainty, to evaluate the effect of GH stability on estimated muscle activity (c).*

Cost function

To find a physically plausible solution $\{\boldsymbol{a}_k, \boldsymbol{c}_k\}$ for each timestep t_k of the experimental motions, we formalize our cost as a convex function to minimize the sum of weighted-squared muscle activations, to represent muscle contribution to the perceived effort:

$$J(\boldsymbol{a}_k, \boldsymbol{c}_k) = \sum_{i=1}^{N_m} w_i a_{i,k}^2 + \sum_{j=1}^{N_\theta} \mu_j c_{j,k}^2 \,, \tag{2.1}$$

where w_i and μ_j are weightings to encourage the use of muscles over reserve actuators, and thus we set them to 1 and 10 respectively. We regularized the range of values of $c_{j,k}$ to be similar to that of the activations $a_{i,k}$, by appropriately selecting the corresponding maximum reserve forces, which are scaled by $c_{j,k}$ when producing ideal torques on each of the joints. In this way, the use of force generated by the ideal reserves is 10 times more costly than that generated by the muscles. This achieves *O1* in terms of solving a muscle redundancy problem, where the simple quadratic formulation is efficient to minimize.

Constraints

Setting the state of the model a priori allows the formulation of simpler and more manageable constraints, which guide the solver towards a feasible solution. Muscles' paths, positions, and inertia of the various bodies are also fixed based on the model's state. Recognizing that the force produced by each muscle is linear in its activation, the previous observations mean that joint accelerations depend linearly on the optimization variables as well. With some manipulation, the simulated accelerations can be

constrained to match the experimental ones by enforcing:

$$\boldsymbol{A}_{\mathrm{acc},k}\begin{bmatrix}\boldsymbol{a}_k\\ \boldsymbol{c}_k\end{bmatrix} = \tilde{\tilde{\ddot{\boldsymbol{\theta}}}}_k\,, \tag{2.2}$$

where $\tilde{\tilde{\ddot{\boldsymbol{\theta}}}}_k \in \mathbb{R}^N$ is obtained by subtracting the accelerations induced by gravity and passive muscle forces from the experimentally recorded accelerations, and the element $A_{\mathrm{acc},k}(j,i)$ of $\boldsymbol{A}_{\mathrm{acc},k} \in \mathbb{R}^{N_\theta \times (N_\mathrm{m}+N_\theta)}$ represents the effect of a unitary activation/control of actuator i on the acceleration of coordinate j. Note that this allows for correctly scaling the contribution of each actuator to the resulting coordinate's acceleration, while the effects of gravity, external forces, and passive muscle forces are lumped in $\tilde{\tilde{\ddot{\boldsymbol{\theta}}}}_k$ (achieving *O3*).

Similarly, we observe that the expression for the JRF at each joint is provided by a vector sum of the moments and forces produced by each actuator. Under the same conditions reported above, we can write the JRF at a generic joint in the model as

$$\boldsymbol{F}_k = \boldsymbol{A}_{\mathrm{F},k}\begin{bmatrix}\boldsymbol{a}_k\\ \boldsymbol{c}_k\end{bmatrix} + \boldsymbol{F}_{0,k}\,, \tag{2.3}$$

where $\boldsymbol{F}_{0,k} \in \mathbb{R}^3$ denotes the value that the reaction force would assume at the current state if all the actuators were de-activated (accounting also for external forces, if present), and the elements of $\boldsymbol{A}_{\mathrm{F},k} \in \mathbb{R}^{3\times(N_\mathrm{m}+N_\theta)}$ account for the effect that unitary activations/controls of the individual actuators would have on the resulting components of the reaction force. Once again, the formalism adopted allows for considering only the effects of the active muscle forces, while the contributions of the passive forces and gravity are lumped into $\boldsymbol{F}_{0,k}$. The stability of the GH joint is then enforced by constraining the direction of JRF at the glenoid to intersect the glenoid fossa, whose shape we approximate as circular (achieving *O2*). Similarly to [97], this results in:

$$\left(\frac{\psi_k(\boldsymbol{a}_k,\boldsymbol{c}_k)}{\psi_{\max}}\right)^2 - 1 \le 0\,, \tag{2.4}$$

where $\psi_{\max} \in \mathbb{R}$ stands for the maximum allowable angle that the reaction force can assume with respect to the line joining the glenoid center and the humeral head, and ψ_k represents the angle produced by the current value of the optimization parameters. This results in a convex constraint, easier to handle in optimization. We chose a conservative value of $\psi_{\max} \approx 20^\circ$, in agreement with cadaveric studies [67].

A final set of constraints was applied to limit muscle activation and extra actuators' controls within the bounds of physiological activation and deactivation rates. The muscle activation in the model represents the level of muscle fiber calcium release resulting from depolarization, where 0 signifies that no calcium is released and the fiber produces no active tension, and 1 means maximum calcium release (activation) and that the fiber contracts maximally. Integrating the activation dynamics in [101] between consecutive time instants t_{k-1} and t_k, we can formulate an expression that enforces activation

dynamics implicitly (achieving *O3*):

$$\underbrace{a_{i,k-1} - a_{i,k-1}\left(\frac{1}{2} + \frac{3}{2}a_{k-1}\right)\frac{t_k - t_{k-1}}{\tau_{\text{deact}}}}_{l_{i,k}(a_{i,k-1})} \leq a_{i,k} \leq \underbrace{a_{i,k-i} + \frac{1 - a_{i,k-1}}{\frac{1}{2} + \frac{3}{2}a_{k-1}}\frac{t_k - t_{k-1}}{\tau_{\text{act}}}}_{u_{i,k}(a_{i,k-1})}, \tag{2.5}$$

written for muscle i, where τ_{act} and τ_{deact} are the activation and deactivation time constants of the muscle, that we set to 10 ms and 40 ms respectivey [101]. The formulation in equation (2.5) bounds the admissible activation's changes by considering the effect of minimum or maximum neural excitation input on the current activation level. The resulting $u_{i,k}(a_{i,k-1})$ and $l_{i,k}(a_{i,k-1})$ should be clipped if necessary, to always respect $a_{i,k} \in [0, 1]$.

2.2.2. Shoulder musculoskeletal model

We employed the thoracoscapular shoulder model [58], which was already scaled to the subject whose movements we will consider [58]. The model (Figure 2.2) was previously used to capture the work done by the primary muscles of the shoulder and includes the kinematics of a 4-DoF scapula relative to the thorax, and the glenohumeral joint as a 3-DoF gimbal joint. The full mobility of the scapula is restricted by the clavicle by the sternoclavicular and acromioclavicular joints. Overall, the shoulder model features 7 DoFs and is actuated by 33 muscle elements.

The previous study investigated the work of shoulder muscles during common shoulder movements, but did not consider the stability of the GH joint. In this study, the role of the rotator cuff muscles (infraspinatus, supraspinatus, teres minor, and subscapularis, highlighted in Figure 2.2) is of special interest since they are anatomically considered glenohumeral joint stabilizers and their activity cannot be measured via surface EMG.

To ensure that the model successfully tracks all the experimental joint trajectories, we included ideal actuators to each coordinate in the model [100]. The estimated muscle activations are realistic if the force/moments that the extra actuators deliver are negligible compared to the muscles. We refer to all the muscles and ideal actuators as "actuators".

2.2.3. Experimental motions and EMG data processing

We considered marker data and surface EMG signals for 11 muscles recorded during 18 experimental trials, at 100 Hz and 1000 Hz respectively, from a previous study [58] where the data is openly accessible. The same subject executed three repetitions of shoulder forward flexion, abduction and shrugging, with and without a 2kg load in hand, resulting in 18 acquisitions divided into 6 tasks. The electrical activity of the following muscles was measured via surface EMG electrodes and processed: anterior, middle, and posterior trapezius and deltoids, pectoralis major, teres major, serratus anterior, latissimus dorsi, and infraspinatus. The marker data corresponds to 9 marker trajectories, representing the motion of several bony landmarks and virtual markers extracted from post-processing of the original motion.

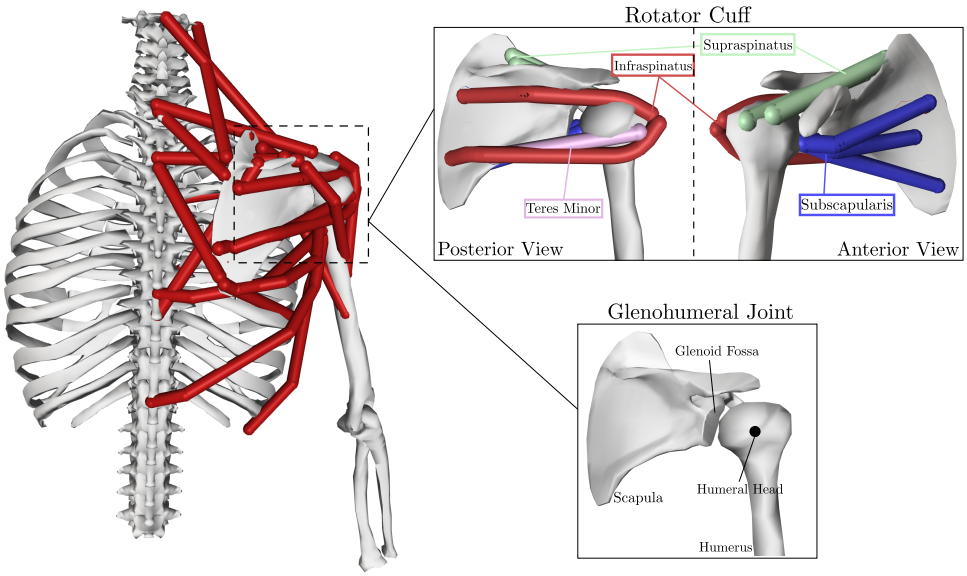

Figure 2.2: *The thoracoscapular shoulder model, with emphasis on the muscles (red lines), the rotator cuff, and structures of the GH joint. In particular, the rotator cuff muscles are highlighted (infraspinatus - red, supraspinatus - green, teres minor - pink, subscapularis - blue).*

From experimental marker trajectories, we calculated joint angles of the shoulder model using the inverse kinematics (IK) tools in OpenSim. We were concerned about reproducing a more consistent ratio between scapula upward rotation and shoulder elevation, as reported in a previous bone pin study [102] and previously experimentally observed ratios [103]. The difference was the introduction of low scapula coordinate weights (2×10^{-4} for the scapula upward rotation, 10^{-4} for the other coordinates) to slightly penalize deviations from its neutral pose. Despite the correction, root mean square error between the experimental and simulated markers' trajectories was below 1 cm on average, across the whole dataset. The resulting joint angles were filtered with a 4^{th} order low-pass Butterworth filter with a cut-off frequency of 3 Hz to eliminate high-frequency noise coming from the marker data (similarly to what was done in [58]), and then differentiated to find the joint velocities and accelerations corresponding to the movement of the subject.

2.2.4. Robustness analysis

To determine the effects of the GH constraint on the individual muscle activations, we took into account the uncertainty of the model fit to the marker data to increase the robustness of our conclusions. Our analysis addresses the issue that for any given marker error, there are infinite solutions for joint angles that can be within that error

and inverse kinematics simply provides one solution, which is not the mean or the most representative of all possible solutions. Consequently, when employing inverse kinematic to estimate the joint angle trajectories corresponding to the marker data, the uncertainty in the placement of the markers or in the scaling of the model can have a significant impact on the conclusions of a biomechanical study [104]. We focus on the uncertainty from marker placement on the model, as the uncertainty of model scaling affected derived joint angles, moments and powers in a very similar way [104].
Starting from our initial model, we generated 100 additional models by perturbing the position of each marker randomly, inside a sphere of 1 cm radius centered on the initial marker position. When generating a model, we ran inverse kinematics on all the marker trajectories in the dataset, and retained the model only if the associated average root-mean-squared error was below 1 cm [104]. We estimated the individual muscle activations of each model and associated inverse kinematics results using the RMR solver and then evaluated if muscle activations with and without the GH constraint enforced were statistically different across 3636 solutions. The 18 experimental trials for both conditions (36 solutions) were perturbed 100 times to ensure our results were robust to uncertainty from IK, for a total of 3636 solutions. To analyse the differences induced by GH stability, we employed the statistical parametric mapping (SPM) method to perform a paired t-test among the sets of one-dimensional time series describing the activations of each muscle, estimated under the two stability conditions. For every muscle, we tested the null hypothesis that the mean activation trajectory in the two conditions is the same, setting the level of significance for the test to be $\alpha = 0.01$, and leveraging the MATLAB interface of the freely available package SPM1D [105] for this analysis. Even if the two conditions achieved a p-value $< \alpha = 0.01$, this may not signify meaningful differences in muscle recruitment, and accordingly, we also considered the effect size for each muscle [106]. The effect size was evaluated as the difference between the means of the activation trajectories between the two conditions. An effect size whose absolute peak value exceeded 0.1 was considered significant since such variation could corrupt significantly its match to experimental EMG values. As such, we identified a muscle to be significantly affected by the GH constraint if the requirements on p-value and effect size are both simultaneously satisfied when comparing its activations under the two stability conditions.

2.2.5. Implementation and simulations

A total of 3636 RMR simulations and 36 CMC simulations were run on a Dell Latitude 7420 laptop with i7-1185G7 processor. The NLP problem equation (2.1)-equation (2.5) addressed by the RMR solver was coded leveraging OpenSim 4.3 in MATLAB R2021b and fed to the numerical solver SQP (available through *fmincon*). To cope with the gimbal lock arising in the model when the humerus is vertical, simulations of shrugging tasks were achieved locking the two indeterminate coordinates (axial rotation and plane of elevation of the humerus), which varied little through the shrugging movement and followed from previous simulations in [58]. The data used in this study is available at https://simtk.org/projects/thoracoscapular, while the code can be accessed at https://github.com/ComputationalBiomechanicsLab/rmr-solver/tree/v1.1.

2.3. Results

We present our results beginning with the accuracy of muscle activations estimated by the RMR solver compared to experimental EMG and to CMC results as a benchmark. Next, we quantify the effect of enforcing GH stability on estimated muscle activations, while taking into account the uncertainty in the movement data.

2.3.1. Comparison between RMR and CMC

Muscle activations estimated by the RMR solver and the CMC algorithm were compared to the filtered EMG signals recorded experimentally. Averaging the activation estimates over the 3 repetitions of each task, we performed a task-wise comparison. Results are reported for the RMR solver with and without the GH constraint enforced since the latter offers a fairer comparison of our solver with respect to CMC, which does not take GH stability into consideration. In Figure 2.3 we show, for a selection of muscles, the mean and standard deviations of muscle activations estimated by the three methods over the duration of the task and specifically for the loaded flexion task. To capture the differences across the whole dataset, we present the mean absolute errors (MAEs) between model-estimated and EMG-recorded muscle activations over the duration of each task, and summarize them in Table 2.1. The values of MAEs show that activations estimated with our method and the popularly employed CMC algorithm are similar, but when visualizing the direction of the JRF at the glenoid computed by the RMR solver, major changes that are induced when including GH stability in the analysis are displayed, as presented in Figure 2.4.

We compared the computation-to-real-time ratio for the algorithms, highlighting significant differences in performance over the whole dataset. The ratio is over 200 for the CMC algorithm, while an average close to 20 was observed for the RMR solver (both in the case in which GH stability was excluded and when it was included), making it a much more efficient option to achieve the estimates. We visualize this comparison in Figure 2.5, including the performances of the RMR solver when the JRF is retrieved through a force-balancing procedure by employing the OpenSim API directly, similar to [98].

2.3.2. Glenohumeral stability analysis

To investigate the effect of the GH stability constraint, we compared the activations estimated by the RMR solver from the two conditions, for all the muscles. We included both the superficial muscles (some of which were monitored experimentally with EMG) and the deeper muscle such as the rotator cuff muscles for which there were no EMG measurements. We evaluated the set of 3636 solutions considering the uncertainty in the movement data to discern significant and meaningful differences in the activation patterns between the GH stability conditions. Figure 2.6 presents the mean and standard deviation of the estimated muscle activations and regions of statistical difference (shaded) and the corresponding effect sizes (below each plot) for the loaded abduction task. We are particularly interested on the effects on the rotator cuff muscles and prime movers whose activation varied the most among the two stability conditions. The muscles in which the peak effect size (the peak absolute difference between the

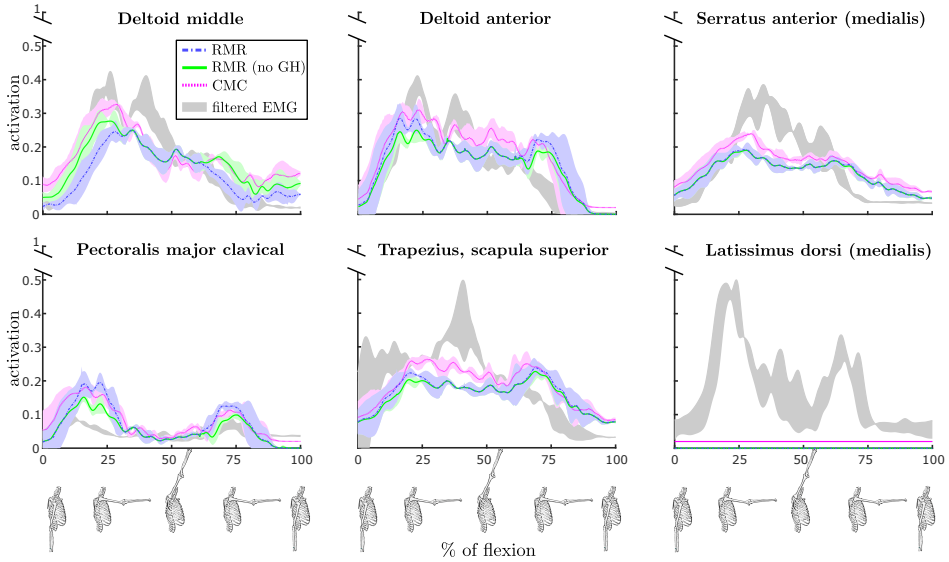

Figure 2.3: *Comparison of muscle activations during a loaded flexion movement of the shoulder as estimated by the RMR solver with (blue) and without (green) GH constraint and by the CMC algorithm (magenta). Additionally, actual muscle activations obtained by filtered measured EMG signals are displayed in gray. For each shown muscle, estimates are displayed on vertical axes in terms of the mean across the 3 repetitions of the task (with bold lines), together with shaded ±1SD, while only the ±1SD region is reported for EMG. The horizontal axes show the progression of the movement (where 0 is the beginning and 100 is the final sample), also visually indicated with skeletal models at the bottom of the figure. The figure provides an intuitive example including the primary muscles to elevate the humerus, while an overview of all the muscles and movements is provided in Table 2.1.*

two means) exceeds 0.1 are starred. In Table 2.2 we present the comparison of all the muscles in the model, summarized by the peak effect size for each task. The shaded values in the table represent the muscles whose activations are identified as physiologically different (p-value < 0.01, peak effect size > 0.1) between the two GH stability conditions.

2.4. Discussion

The main aim of our study was to implement and analyze the effect of enforcing glenohumeral stability on the estimation of muscle activations during human shoulder movements. To that end, we developed a new open-source Rapid Muscle Redundancy (RMR) solver compatible with OpenSim in order to capture the GH stability constraint, include passive muscle fiber force contributions, and limit the physiological rate of change of estimated activations for individual muscles.

Task	Trapezius middle	Trapezius superior	Trapezius inferior	Deltoid anterior	Deltoid posterior	Deltoid middle	Pec.Maj. clavicle	Serratus anterior	Infra-spinatus	Latiss. dorsi	Teres major	Method
Abduction	0.09	**0.11**	0.06	0.09	0.07	0.07	0.02	0.07	0.08	0.06	0.06	RMR
	0.09	0.10	0.06	0.09	0.07	0.08	0.02	0.07	0.05	0.06	0.03	RMR (no GH)
	0.07	0.10	0.05	0.08	0.07	0.07	0.01	0.06	0.06	0.04	0.02	CMC
Abduction +	0.07	**0.11**	0.05	0.07	0.10	0.04	0.03	0.06	0.05	**0.13**	0.08	RMR
	0.07	**0.11**	0.05	0.06	0.09	0.06	0.03	0.06	0.09	**0.13**	0.05	RMR (no GH)
	0.06	0.10	0.04	0.05	0.09	0.07	0.02	0.07	0.08	**0.11**	0.04	CMC
Flexion	0.06	0.07	0.04	0.05	0.05	0.05	0.02	0.05	0.04	0.10	0.07	RMR
	0.06	0.07	0.03	0.05	0.05	0.05	0.02	0.05	0.04	0.10	0.06	RMR (no GH)
	0.05	0.06	0.03	0.05	0.03	0.05	0.03	0.05	0.02	0.08	0.05	CMC
Flexion +	0.04	0.08	0.02	0.05	0.07	0.04	0.04	0.06	0.02	**0.16**	**0.11**	RMR
	0.04	0.08	0.02	0.04	0.07	0.05	0.02	0.06	0.04	**0.16**	0.10	RMR (no GH)
	0.03	0.06	0.02	0.05	0.06	0.06	0.04	0.06	0.02	**0.15**	0.09	CMC
Shrugging	0.02	0.09	0.01	0.01	0.02	0.02	0.06	0.01	0.08	0.03	**0.17**	RMR
	0.02	0.10	0.01	0.00	0.02	0.02	0.05	0.01	0.08	0.03	0.10	RMR (no GH)
	0.01	0.08	0.01	0.01	0.01	0.01	0.05	0.04	0.06	0.02	0.01	CMC
Shrugging +	0.04	0.06	0.02	0.01	0.02	0.02	0.05	0.02	0.02	0.10	**0.12**	RMR
	0.04	0.07	0.02	0.00	0.02	0.02	0.05	0.01	0.02	0.10	0.03	RMR (no GH)
	0.02	0.05	0.01	0.01	0.02	0.06	0.01	0.07	0.01	0.08	0.02	CMC

Table 2.1: *Mean Absolute Error (MAE) for RMR and CMC estimates of the muscle activations against the EMG-based activations. Results for the RMR solver are shown for both the cases in which the GH constraint was included or not. MAEs under 0.1 represent excellent matches with the experimental activations, while worse values are highlighted. For the serratus anterior, both RMR and CMC estimates are averaged over the 3 bundles composing the muscle in the model, as reported in [58].*

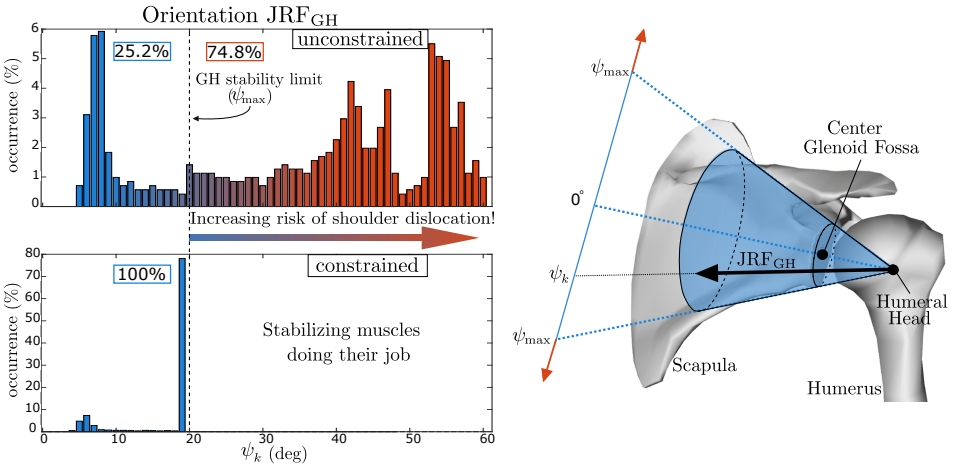

Figure 2.4: *The effects of the constrained joint reaction force (JRF) at the glenoid during a loaded forward flexion (left) and the physical interpretation of the glenohumeral (GH) stability constraint (right). Since the JRF physiologically must remain within the glenoid fossa, our algorithm constrains its orientation ψ accordingly. The two histograms present the values of ψ in the two cases, reporting the rate (%) at which angles occur during the movement. Without ensuring GH stability, it is evident that the glenohumeral JRF can grossly exceed joint stability limits.*

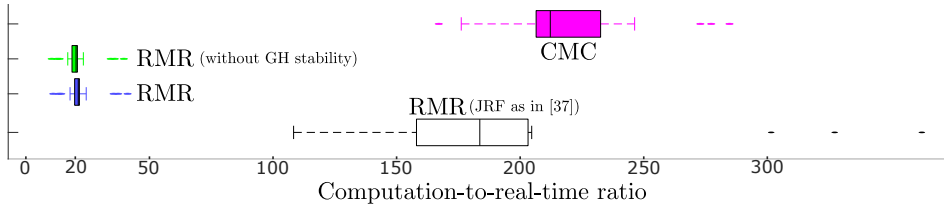

Figure 2.5: *Comparison of computational performance of the RMR solver versus the CMC algorithm, over the 18 experimental motions in our dataset collected at 100 Hz. For the RMR solver, both the cases with and without the inclusion of the glenohumeral stability constraint are included and both indicate a processing rate of nearly 5 frames per second. For comparison, we include the RMR Solver formulation that includes the JRF computed from the multibody system at every constraint evaluation [98] instead of a linear function of activation at each instant.*

Muscle	Flexion	Flexion+	Abduction	Abduction+	Shrugging	Shrugging+
Infraspinatus	0.07	**0.14**	**0.10**	**0.17**	0.02	0.06
Supraspinatus	0.04	0.09	0.03	0.07	0.04	0.03
Subscapularis	0.03	**0.13**	**0.17**	**0.27**	0.07	0.08
Teres minor	0.02	0.05	0.04	0.05	0.02	0.04
Trapezius	0.02	0.03	0.01	0.01	0.03	0.03
Deltoid anterior	0.03	0.05	0.02	0.07	0.01	0.01
Deltoid middle	0.03	0.08	0.03	0.05	0.00	0.00
Deltoid posterior	0.01	0.06	0.04	0.05	0.00	0.01
Pectoralis major	0.03	0.07	0.01	0.02	0.02	0.01
Serratus anterior	0.00	0.01	0.00	0.01	0.04	0.04
Latissimus dorsi	0.00	0.00	0.00	0.00	0.00	0.00
Teres major	0.04	**0.10**	**0.12**	**0.12**	**0.13**	**0.11**
Levator scapulae	0.00	0.00	0.00	0.00	0.04	0.02
Coracobrachialis	0.02	0.04	0.01	0.03	0.02	0.02
Pectoralis minor	0.00	0.00	0.00	0.00	0.00	0.00

Table 2.2: *Peak effect size between the with and without GH stability conditions for estimated muscle activations for all shoulder muscles modeled. Shaded values indicate that the peak effect size* ≥ 0.1 *which corresponds to a meaningful difference in muscle coordination. For the multi-bundle muscles, we show the maximum peak effect size across them.*

2.4.1. Effect of GH stability on muscle activations

With respect to the effect of enforcing GH stability on muscle activations of the shoulder, our result showed that GH stability increased rotator cuff muscle activity (Figure 2.6, Table 2.2). Furthermore, the increase in muscle activity when accounting for GH stability is greater when the same movements are performed with a handheld weight. Interestingly, GH stability had virtually no effect on activations of surface muscles, whose EMG recordings are typically used to validate estimated muscle activity [58, 80, 86, 95, 107, 108]. Indeed, our comparisons indicated the RMR estimates with and without GH

stability were comparable to CMC with respect to filtered EMG signals, with all methods having excellent agreement (MAE ≤ 0.1, Table 2.1).

Our findings are consistent with anatomical expectations that the rotator cuff muscles act to stabilize the glenohumeral joint [65–67]. Previous shoulder modeling studies also indicated the importance of the rotator cuff muscles in controlling both magnitude and direction of the glenohumeral joint reaction force [89, 92, 96]. Yet, there are a couple of studies that observed that constraining the GH joint reaction force did not significantly affect the activations of the rotator cuff muscles [95] or observed marginal differences [98]. The models in these studies and our model, however, represent different individuals with different geometry and muscle architecture (compare Holzbaur et al. [109] to van der Helm et al. [89, 110]) and employ completely different models of scapula kinematics. We specifically model the scapulothoracic joint [102], which was attributed with improving estimates of the work done by individual shoulder muscles [58]. Differences between our studies extend to different methods to solve for the muscle redundancy, particularly in the study by Blanche et al. [98], where they employed static optimization. Nonetheless, our study supports the recommendations to include a constraint on the JRF at the glenoid when investigating GH stability [95].

2.4.2. Computational speed of RMR solver to estimate forces and activations

In formulating the RMR solver, the secondary aim was to improve the computational efficiency of estimating muscle forces while maintaining (or even improving) the accuracy of methods like CMC [70]. Our results indicate that the RMR solver is over an order of magnitude more efficient at estimating muscle activations in a complex musculoskeletal model of the human shoulder than the widely employed CMC algorithm.

The speed gains of the RMR solver are attributable beyond the efficiencies of static optimization. In particular, it is a result of a novel formulation of the JRF as a linear function of activation within the constraints of the RMR solver as opposed to repeatedly querying the multibody system for reaction forces in response to muscle forces [98] (Figure 2.5). In the case of the shoulder, we demonstrated the efficient inclusion of JRFs by implementing the directional constraint on the glenohumeral joint reaction force.

Notably, our JRF formulation can be employed to include the reaction forces at any joint (in any OpenSim model), for example, to keep JRFs within physically acceptable bounds or to append them to the objective function to be minimized. The efficient inclusion of JRFs can be used to improve the physiological feasibility of estimated muscle forces.

The RMR solver could solve the muscle redundancy problem formulated in equations (2.1) and (2.5) in approximately 0.2 seconds for a single instant in time, and directly enables motion analysis at up to 5 frames-per-second. These reduced computational times pave the way for musculoskeletal model-based analysis of human motion in near real-time, which opens opportunities to apply modeling insights during rehabilitation and therapy. Our study exploited the improved computational efficiency, by enabling us to perform a robustness analysis with over 3600 trials to test against the uncertainty in the motion data. Such large-scale robustness, sensitivity and even design optimization problems would be impossible to perform in tractable time with CMC.

2.4.3. RMR SOLVER ACCURACY OF MUSCLE ACTIVATION ESTIMATES

Estimated muscle activity from the RMR solver compared to the experimentally recorded EMG signals and, with few exceptions, the MAE observed between the two remained below the threshold of 0.1, commonly accepted as indicating excellent agreement between simulated and real muscle activation levels (see [58, 86, 111]). Reserve actuator torques were negligible for the whole duration of the movements, confirming the relative weightings selected in the cost function equation (2.1). Our results indicate that the integration of the system's dynamics performed in CMC, together with its kinematic feedback loop, may not be necessary when analyzing relatively slow movements, like the ones investigated in this study, and supports previous studies comparing static to dynamic optimization methods for estimating muscle forces in walking [112]. Our approach further supports the observation that a rigid tendon model, disregarding tendon compliance, performs just as well for a wide variety of shoulder motions [113].

2.4.4. LIMITATIONS

Our results are subject to several limitations. We have modelled the GH constraint as a directional constraint, on the basis of a circular approximation of the glenoid fossa. However, this may underestimate the support provided by the actual shape of an individual's glenoid fossa and similarly by excluding soft-tissue structures such as the glenoid labrum, joint capsule, and glenohumeral ligaments in the simulations. The constraint on JRF direction in equation (2.4) results in orientations of the reaction force that lie predominantly on the boundaries of the admissible region (e.g., the perimeter of the glenoid fossa), as apparent in Figure 2.4. Other orientations of the JRF could be encouraged as well, but it is important to note that they would likely require higher activations than the ones we estimated for the rotator cuff muscles. Therefore, the muscle activation estimates that we report may be seen as the lower bound required for ensuring GH stability, and beyond this lower bound, we would expect increased influence of GH stability on the rotator cuff muscle activations. The approximation of the acceptable region based on the glenoid fossa and the subsequent directional constraint may also explain some of the outliers in our results. In particular, teres major (Table 2.2 and Table 2.1) appears to be overestimated particularly during shrugging with GH stability constraints active. We believe this may be due to its role in redirecting the forces from the trapezius into the GH fossa, even though the magnitude of the force is small, which other soft tissue like the capsule and ligaments could support.

Another outlier in all cases (using CMC and RMR) was the latissimus dorsi (see Table 2.1), particularly during arm-raising tasks with a handheld weight, where the model underestimates muscle activity. We do not attribute these differences directly to GH stability since these discrepancies appear in both cases (with and without GH stability). Functionally, we consider latissimus dorsi as an antagonist to the deltoids and superior trapezius, and its activity when the humerus is horizontal indicates a stabilizing role in the experimental EMG data. As the model does not predict latissismus dorsi coactivation (both employing the RMR solver and the CMC algorithm, see Table 2.1) and instead appears to rely on particularly high activations for the teres major to stabilize the joint, we may infer that the muscle path or architecture of the latissmuss dorsi may

not be captured adequately. While this is not specific to the solver, it does highlight an important point of improvement for the model, so that a proper distribution of the necessary forces could be realized.
We accounted for the uncertainty in the motion capture data, which was demonstrated to have a significant effect on the conclusions drawn in human movement studies [104]. We used the marker placement uncertainty to also account for uncertainty in model scaling. However, there is also uncertainty related to the muscle parameters and muscle architecture of the test subject, which we did not account for. Previous work concluded that estimated muscle forces and activations are also sensitive to these parameters, but personalization of the model alone does not ensure high accuracy, while properly addressing the muscle force sharing problem (as we aimed to do) is equally important [114].

2.5. Conclusion

We showed that glenohumeral stability has a significant and large effect on the activations of the rotator cuff muscles. Our results highlight that simulation studies performed without glenohumeral stability can yield non-physical directions for the glenohumeral joint reaction force. Furthermore, the activations of the rotator cuff and teres major muscles are the most affected by the stability constraint. Consequently, we demonstrate that validation of shoulder model estimated muscle activity, by comparison to superficial EMG data, is insufficient since rotator-cuff muscle activity is unaccounted for and superficial muscles are widely unaffected by glenohumeral stability. Our results clearly show that good agreement with superficial EMG can be achieved with significantly different rotator-cuff muscle activations.
Our study has broader implications for musculoskeletal models of the shoulder and related solvers for clinical and rehabilitation applications. For example, shoulder surgery [109], robotic-led shoulder physiotherapy [61, 115], and ergonomics during manual work [116, 117] all rely on the accuracy and speed of force/activity estimates from a musculoskeletal model. Our study contributes significantly to these types of applications by providing free, open-source models and tools to accurately evaluate shoulder function in close to real-time.

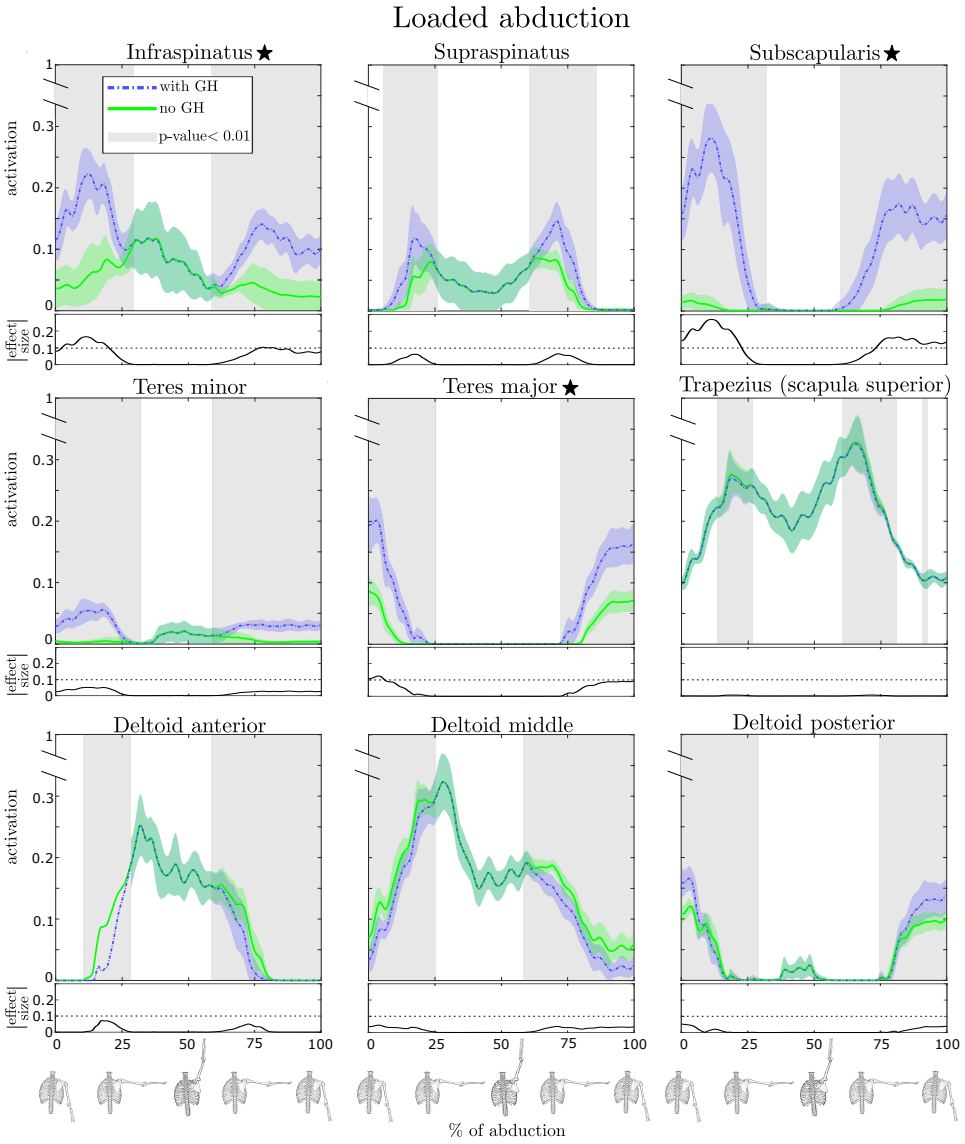

Figure 2.6: *Comparison of estimated muscle activations during a loaded abduction obtained from the RMR solver with (blue) and without (green) GH stability enforced for several key muscles. The vertical axes of the main graphs display muscle activation where value 1.0 is the maximum possible activation. Solid lines in the plots correspond to the activation mean while the shaded regions indicate* ±1*SD. The horizontal axes describe the phase of the movement where 0 is the beginning and 100 is the final sample, which are visually indicated by skeletal poses. The shaded grey sections of the phase indicate where the difference is statistically significant (*$p < 0.01$*). The absolute value of the effect size is shown below the main graph, and we highlight with the muscles for which this curve exceeds* 0.1*. The loaded abduction task (shown) exemplified the differences between the two conditions, but an overview of all muscles for each task is provided in Table 2.2.*

3

Biomechanics-aware robotic navigation of shoulder tendon strains

The previous chapter introduced the RMR solver, a fast algorithm for estimating muscle activations within a musculoskeletal model, and applied it to the human shoulder to investigate deep muscle activity during upper-limb movement. We showed that rotator cuff activity cannot be inferred from superficial muscle signals alone, underscoring the need for musculoskeletal simulations to monitor the level of engagement of these deep muscles. Once muscle activations are known, other related quantities, such as tissue loading, can be estimated too.

Building on this foundation, the present chapter integrates high-fidelity simulations into the real-time control of a rehabilitation robot, aiming to provide patients with personalized, low-tendon-strain shoulder movements during rotator cuff recovery. We propose a Biomechanics-Aware Trajectory Optimization for Navigating dynamically changing maps of rotator cuff tendon strain (BATON), which leverages a reformulation of the RMR solver to account for human-robot interaction forces during the estimation and retrieves rotator cuff tendon strain online. This tight coupling of musculoskeletal modeling and robot control enables the system to adapt to voluntary or reflexive human actions during physical human–robot interaction (pHRI), thereby reducing re-injury risk.

We evaluate BATON in simulation and real pHRI experiments, comparing it to prior state-of-the-art approaches for robotic-assisted shoulder rehabilitation. Results show that BATON achieves lower tendon strain—linked to reduced re-injury risk—while also decreasing peak accelerations and interaction forces, paving the way for broader integration of high-fidelity musculoskeletal models in rehabilitation robotics.

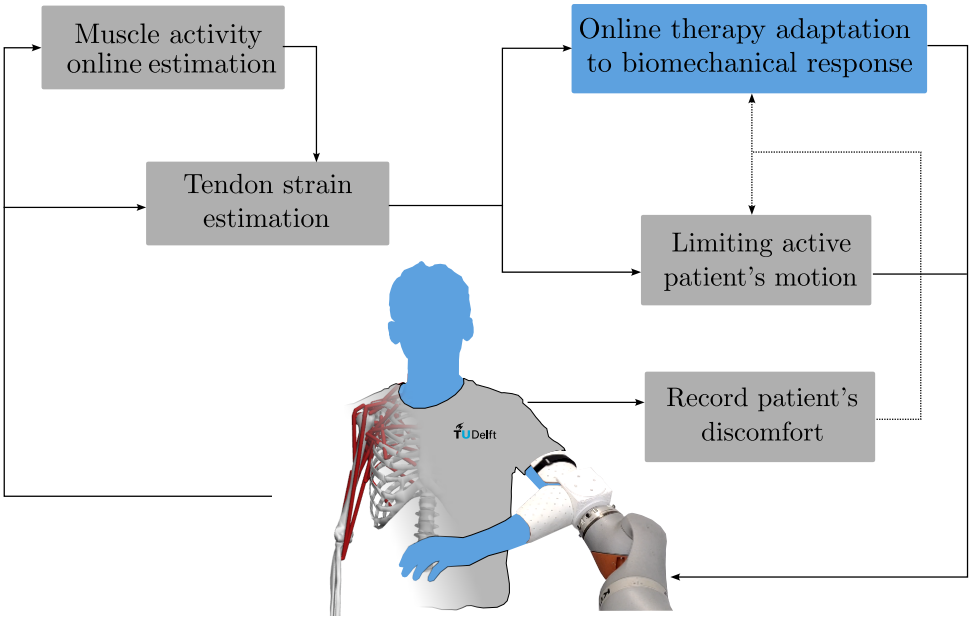

This chapter is based on **I. Belli**, F. van Melis, J. M. Prendergast, A. Seth, and L. Peternel, Biomechanics-Aware Trajectory Optimization for Navigation during Robotic Physiotherapy. Under review, available as preprint at arXiv:2411.03873 (2024). [60].

3.1. Introduction

Injuries affecting the musculoskeletal system are pervasive in society due to an increasingly aging population, strenuous manual labor, and more widespread engagement in sports [48, 118, 119]. Disorders impacting the shoulder and the rotator cuff are particularly common, with clinical literature reporting their prevalence to be as high as 22% in the general population [48] and 50% over age 66 [120]. Overall, the demand for therapy and rehabilitation is large and expected to grow, exacerbating the problem posed by a lack of physical therapists [121]. Moreover, the task of treating a complex biological mechanism, such as the shoulder, is challenging due to a fundamental lack of insights into rotator cuff behavior during rehabilitation. Physiotherapists need to move the joint safely, gradually increasing the patient's range of motion while physically supporting the patient's arm and avoiding re-injuries [14]. Consequently, simple and limited movements are prescribed, to limit both the risk of re-injury and the physical demands on the therapist to manually manipulate the patient safely. However, overly conservative movements may result in a limited range of motion during the therapy, which can delay recovery since moving through a larger range of motion was shown to expedite recovery [122, 123]

Robotic technologies can assist in addressing manual manipulation and re-injury risks during therapy: robots can reduce the manual burden on physiotherapists and provide new tools to monitor and improve rehabilitation outcomes. Overall, rehabilitative robots are already used successfully in post-stroke rehabilitation [20, 124, 125] and gait assistance [21, 126, 127]. However, robots to treat musculoskeletal injuries (such as rotator cuff tears) remain limited. In particular, previous work remains focused on relieving physiotherapists from physically manipulating the patient. Smooth rehabilitation trajectories can be guaranteed by focusing on the control of the movement of the rehabilitative robot itself [128], and therapeutic robotic movements can be learned from expert human demonstrations to automate manual patient manipulation [129]. However, these solutions do not directly account for the patient's biomechanics. Deeper insights into human musculoskeletal mechanics can enable robots to be aware of tissue and joint loads associated with injury risks, allowing them to operate autonomously or in collaboration with therapists to holistically improve the rehabilitation process [130]. An increasingly promising strategy is to leverage human musculoskeletal models and incorporate them directly into robot control loops [40, 41, 131–135].

Musculoskeletal modeling has progressed greatly over the past decades, allowing researchers to estimate the activities of individual muscles involved in producing movement [33, 59, 136] and to perform predictive simulations of human motion [32, 137]. As computational models increasingly capture the inner workings of the human body, their applicability continues to expand, particularly in physical human-robot interaction and assistive devices. Recently, biomechanical models have been used to quantify the assistance needed by a human operator [40, 131, 132], reduce human metabolic cost in walking-assistive devices [41, 133, 134], and even plan search-and-rescue robotic operations [138] or the motion of supernumerary robotics limbs [139]. Other researchers designed optimized trajectories for controlling an ankle rehabilitation robot to minimize joint loading [140]. A custom ankle model was used to plan the robot's trajectory offline, making online model-free adjustments based on tracking errors attributed to excessive

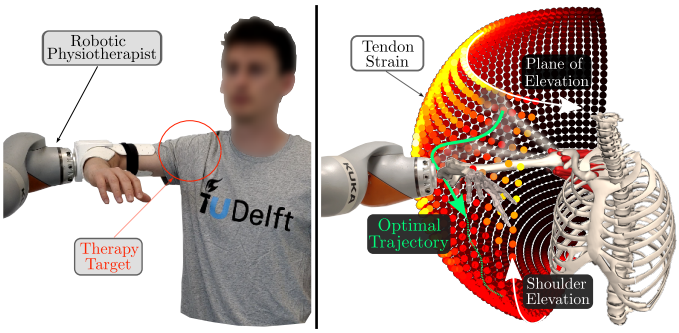

Figure 3.1: *BATON combines musculoskeletal modeling and human-robot interaction to enable a robotic physiotherapist to plan therapeutic movements for its patients in real-time while considering tissue loading (i.e., strain) induced in the rotator cuff tendons in the shoulder.*

joint stiffness. Through computer simulations or offline usage of human models, these works demonstrated the importance and utility of including human biomechanics in robot planning and decision-making. However, they did not leverage high-fidelity biomechanical simulations to regulate human-robot interactions in scenarios that require online adaptation to human behavior, such as robotic physiotherapy.

When real-time capabilities are necessary, purely kinematic models have been used to monitor human joint positions, velocities, and torques [141], or to develop controllers for hybrid neuroprostheses [142]. While lower-fidelity models allow for computational efficiency, they are unsuitable for deep insights into the musculoskeletal system during physical human-robot interaction.

Recently, biomechanical quantities from a high-fidelity human model have been used in an online model predictive control framework, achieving predictive control of a leg prosthesis for ergonomically safe walking [143]. Researchers leveraged fast predictions of how knee torque relates to ankle angle, by including offline information from a musculoskeletal model in an imitation learning strategy. Moya-Esteban and Sartori, instead, directly leveraged surface EMG-driven human models to minimize lumbar joint compression during external load handling, modulating the assistance of a wearable back exoskeleton online [45, 47].

However, these studies did not consider the mechanical behavior of the human tissues (e.g., muscles and tendons) during physical human-robot interaction for shoulder injuries. Without such consideration, safety cannot be ensured throughout the rehabilitation movements. Furthermore, surface EMG is limited to the measurement of surface muscles, while the underlying muscles and tissues cannot be monitored. Since typical shoulder injuries involve the rotator cuff muscles, which are deep-lying muscles that hold the arm in place, there is a need for insights into the underlying biological tissues. To gain such underlying insights, tissue loads can be estimated from external measurements based on a high-fidelity model of the human shoulder. In this direction, researchers developed the concept of “strain maps” [144] to intuitively represent musculoskeletal tissue loadings. A strain map captures the relationship between

human pose and the strain (i.e., load) induced in the rotator cuff tendons. By abstracting the rotator cuff tendon strains into an intuitive and navigable map, offline graph-based motion planning for safe robotic-mediated shoulder rehabilitation was achieved. However, the planning algorithm did not account for the dynamics of the human, resulting in trajectories that presented sudden direction changes, potentially difficult to track, and dynamically unsafe.
Moreover, the method in [144] did not consider online variations in tendon strain level caused by muscle activations during therapy, limiting its real-world applicability. Employing the strain maps, a reactive impedance-control-based approach to physical human-robot interaction in rehabilitation was also presented, to enable a subject to perform therapeutic exercises while robotic assistance protected them from navigating through potentially dangerous poses [61]. In this case, the system would react only when dangerous movements were already initiated, without anticipating the subject's trajectory, which could lead to abrupt movements and high corrective forces from the robot. Critically, unpredictable volitional actions of the patient and/or reflexive reactions to physical human-robot interaction in terms of muscle activity and their effects on tendon strains were not considered.
It is clear that musculoskeletal models have the potential to inform both the planning of the therapeutic motion, in regard to tissue loading, and the human activity and response to robot movement through physical human-robot interaction. However, there has been no single framework for combining planning, execution, and human response during robotic rehabilitation in an online manner.
To address this gap, our overarching objective is to develop a Biomechanics-Aware Trajectory Optimization for the Navigation of real-time human tissue strains during robotic physiotherapy (BATON) targeting musculoskeletal rehabilitation. Our novel approach directly couples a state-of-the-art biomechanical model of the human shoulder to robotic control, generating movements for predictive robotic-assisted musculoskeletal rehabilitation in real-time. This is achieved by accounting for changes in muscle activity due to unpredictable volitional actions of the patient and/or reflexive reaction to physical human-robot interaction that is estimated directly from robot sensors. The resulting planned trajectories are realized by a collaborative robot arm that delivers the movement to a healthy human subject through safe impedance control and provides estimations of contact forces and human poses to close the loop. We propose a muscle activity estimation that provides an alternative to sensor-based muscle activity estimation from surface EMG. Our system can inform the online path planning about the activity of deep-lying muscles, such as rotator cuffs, which cannot be measured with surface EMG.
To achieve our objective, we target the following specific aims:

- **A1**: design a biomechanics-aware trajectory optimization planner that efficiently incorporates a real-time shoulder musculoskeletal model into the underlying optimal control problem (OCP) by decoupling skeletal dynamics and tendon strain behavior of the rotator cuff muscles;

- **A2**: develop a method to inform the planner of the dynamic changes in human biomechanics dependent on muscle activity during real-time physical human-robot interaction;

- **A3**: experimentally evaluate the performance of BATON in terms of the speed, smoothness, and forces of human-in-the-loop response.

Our method, BATON (see Figure 3.2, for overview), extracts insights from the musculoskeletal model for real-time use in optimization and control (Section 3.2.1). To exercise the range of motion of the human shoulder during therapy, the planner finds the minimum strain path by solving a trajectory optimization problem (Section 3.2.2). However, the physical interaction between the robot and the human during trajectory execution induces loads that the human responds to voluntarily and/or reflexively. Since human decisions, sensitivities, and responses to movement and other stimuli are inherently unpredictable, it is imperative that the human response in the control loop is based on actual measurements. Such human-in-the-loop actions and reactions change the level of muscle activity. Thus, muscle forces and corresponding tissue loads cannot be fully predicted, but need to be incorporated. To this end, we estimate these actions and reactions online using a state-of-the-art muscle redundancy solver (Section 3.3.1). We employ a collaborative robotic arm instrumented with a force-torque sensor to estimate the current human pose and loading (Section 3.3.2) and administer rehabilitation movements to the human subject. We embody the robot controller in a closed-loop with the subject (Section 3.4.1) to perform physical experiments (Section 3.4.2) that evaluate our improvements over the state-of-the-art rotator cuff rehabilitation techniques.

3.2. Motion Planning on Strain Maps

Musculoskeletal insights are derived from a human biomechanical model, becoming the strain maps navigated by our online trajectory planner. We developed a rehabilitation path planner to traverse these maps to synthesize safe robotic-led human movements by formulating and efficiently solving an optimal control problem.

To give a practical analogy of how we navigate safety during rehabilitation, let us consider the path planning of a ship in a bay where underwater reefs are present. To know which maneuvers allow the ship to navigate safely and efficiently toward a target position, we need both a dynamic model of the vessel and a map of the underwater reefs, providing non-obvious insights about the safest path. In our situation, the human shoulder dynamics play the role of the vessel, while the strain maps offer insights into what happens beneath the surface (in the tissue being rehabilitated). Safe navigation can only be obtained by considering where it is safe to travel, together with the maneuverability of the ship.

3.2.1. Human Biomechanical Model

We employ a high-fidelity biomechanical model of the human shoulder [58], developed in OpenSim [31, 99], to capture unpredictable volitional actions of the human and/or responses to motions and forces induced by robot interactions. The human model captures the shoulder's movement degrees of freedom (DoFs) and the muscles that generate movement. The model, therefore, can inform the robot's controller of the internal human mechanics during the physical human-robot interaction. In particular,

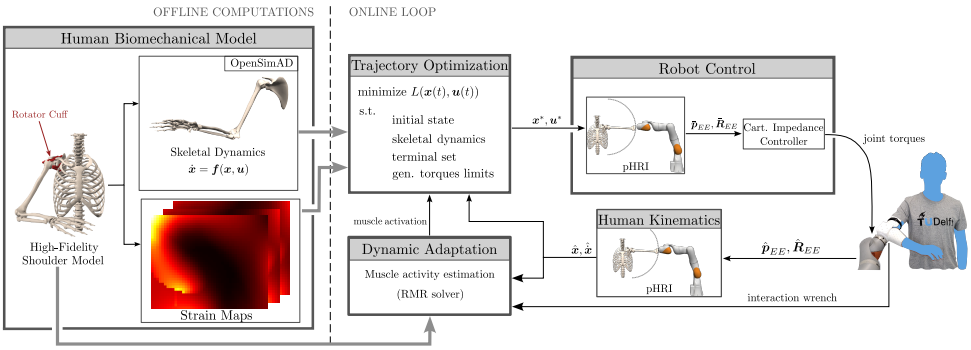

Figure 3.2: *Methodological overview of BATON applied to physiotherapy for the rotator cuff muscle tendons of the human shoulder. The high-fidelity biomechanical model of the human shoulder (left), employed fully to estimate human muscle activations, is decoupled in skeletal dynamics and rotator cuff tendon loading information (strain maps) for efficient trajectory planning. An optimal control problem computes safe rehabilitation trajectories leveraging insights from the human model, and its output is transformed into equivalent references for the robotic controllers to track during physical human-robot interaction (pHRI). The robot's feedback on the current human status is used to close the loop, allowing real-time estimation of the physiological activation for the muscles of interest.*

we are interested in quantities related to injury risk of the rotator cuff muscles (i.e., supraspinatus, subscapularis, infraspinatus, and teres minor), which act as stabilizers for the glenohumeral (GH) joint in the shoulder and whose tendons are subject to tears due to overload, chronic weakness, and impingement. With the model, we predict how movements and forces affect tendon strain in the rotator cuffs, which cannot be measured by surface EMG due to their position deep inside the body. Mechanical strain is associated with re-injury risk of rotator cuff tendons, as excessive stretching of a healing tendon could lead to tearing [145].

In the human body, the rotator cuff spans the GH joint, which permits mobility of the upper arm (humerus) with respect to the shoulder blade (scapula) through 3 DoFs. To improve efficiency without sacrificing fidelity, we reduced the original model to the 3 DoFs and 22 muscles that actuate the GH joint only, to isolate the effects on the rotator cuff muscles (see Figure 3.2, left). The arm dimensions and mass of our human subject were used to personalize the model's properties through OpenSim's scaling functionalities and tabulated anthropometric data [146]. We define the state of the resulting human shoulder model as:

$$\boldsymbol{\theta}_{\mathrm{S}} = [PE, SE, AR]^T \tag{3.1}$$

$$\boldsymbol{a} = [a_1, a_2, ..., a_{22}]^T, \; a_i \in [0, 1] \; \forall i \tag{3.2}$$

$$\mathbf{x} = [\boldsymbol{\theta}_{\mathrm{S}}^T, \dot{\boldsymbol{\theta}}_{\mathrm{S}}^T, \boldsymbol{a}^T]^T \tag{3.3}$$

where PE, SE, and AR are the 3 movement DoFs defined as the Y-X'-Y" sequence of intrinsic Euler angles in the fixed shoulder reference frame (PE: plane of elevation, SE: shoulder elevation as rotation around -X, AR: axial rotation, visualized in Figure 3.3 and indicated collectively as $\boldsymbol{\theta}_S \in \mathbb{R}^3$), $\dot{\boldsymbol{\theta}}_S \in \mathbb{R}^3$ is the vector containing their derivatives with respect to time, and $\boldsymbol{a}$ contains the muscle activation level for each of the muscles in the model. Note that, in the remainder of this paper, we will refer to $\mathbf{q} = [\boldsymbol{\theta}_S^T, \dot{\boldsymbol{\theta}}_S^T]^T$ as the purely kinematic state of the model, while $\mathbf{x}$ is the complete musculoskeletal model state. The personalized shoulder model specifies how the shoulder moves under the influence of externally applied torques/forces, and how these affect the underlying musculoskeletal system. However, embedding the model directly to plan rehabilitation movements through the OpenSim API is not computationally tractable for real-time applications. To achieve real-time performances without sacrificing model accuracy, we decouple tissue mechanics that describe tendon strain with shoulder movement and muscle activation from skeletal dynamics that simulate how passive shoulder motion is generated as a consequence of externally applied forces. This decoupling enables us to formulate the problem of navigating shoulder motion over strain maps governed by skeletal dynamics and to solve planning efficiently.

Navigable strain maps

A strain map describes how strains in individual tendons, or collectively (e.g., the peak), change with respect to the model state (motion and activation). As the complete human model state $\mathbf{x}$ spans a high-dimensional space, we focus on the pose (target position from the therapist) and overall level of activation as the primary dimensions of a spatial strain map. We compute the strain map querying the musculoskeletal model to extract quantitative information on the mechanical strain σ_i of every i^{th} tendon in the rotator cuff. Biologically, tendon strain depends on the model's state $\mathbf{x}$ as:

$$\sigma_i = \frac{l_i(\mathbf{x}) - l_{i,0}}{l_{i,0}} \cdot 100\%. \tag{3.4}$$

where $l_i(\mathbf{x})$ is the current length of the tendon and $l_{0,i}$ its slack length at rest. We fully describe the change in strain σ_i based on the human model's state via pre-computed "strain maps" [61, 144, 147], for every tendon of interest. In this way, we created high-dimensional strain maps for every individual tendon, which allowed us to target only one tendon at a time or, alternatively, to monitor the strain on the rotator cuff as a whole by constructing aggregate maps that consider the highest strain across all the tendons [144, 147]. To simplify the understanding of strain level for non-expert users of our system (e.g., physiotherapists and patients), we chart σ_i on 2D grids where the strain percentage is visualized as a function of PE and SE angles, while both AR and a_i are fixed. To consider different values of AR and a_i, other (pre-computed) 2D maps can be traversed.

To traverse over strain maps efficiently, we interpolated our discrete strain maps using 2D Gaussian radial basis functions. Through least-squares fitting, we obtain the parameter vector $\boldsymbol{\pi} \in \mathbb{R}^{N_p \times N_G}$ that describes the combination of the N_G Gaussian

functions best approximating every 2D strain map, in terms of the N_p parameters of every Gaussian.

Controlling skeletal dynamics

After computing navigable differentiable maps of rotator cuff tendon strain, we account for the human movement dynamics when the human shoulder is being manipulated by the robot. Human dynamics should be used to inform the trajectory planner (Section 3.2.2) about the relationship between forces/torques that the robot applies to the human, and the resulting human accelerations. However, the OpenSim API provides this relationship only in a non-differentiable form, which led previous research [139] to employ less efficient gradient-free optimization methods to optimize human-robot interaction. To efficiently represent human skeletal dynamics for control purposes, we employed OpenSimAD [148] as a tool that generates differentiable functions from OpenSim musculoskeletal models. OpenSimAD extends OpenSim by recording the call sequence used to compute specific outcome variables within an OpenSim model, enabling algorithmic differentiation through frameworks like CasADi [149]. In particular, we customized the OpenSimAD pipeline to obtain a differentiable representation of our human model in terms of the ordinary differential equations capturing the dynamics of its multibody skeletal system, given a set of generalized forces $\mathbf{u} \in \mathbb{R}^N$ and the model's kinematic state $\mathbf{q} \in \mathbb{R}^{N_q}$:

$$\dot{\mathbf{q}} = f(\mathbf{q}, \mathbf{u}), \tag{3.5}$$

where $N = 3$ and $N_q = 2N$ in our case. Importantly, this procedure guarantees that the model's joint definitions will be preserved, together with personalized dimensions, mass, and inertia for every existing bone in the original model.

3.2.2. Biomechanics-aware trajectory optimization

We combine information from the differentiable strain maps and human skeletal dynamics to formulate a biomechanics-aware optimal control problem (OCP), with the goal of synthesizing safe human motion towards a target shoulder position. The solution to the OCP is the optimal trajectory $\mathbf{q}^*$ and corresponding generalized forces $\mathbf{u}^*$ driving the human model toward the goal pose. Note that we optimize generalized forces rather than end-effector wrenches to keep our trajectory optimization agnostic to how rehabilitation trajectories are delivered to the patient. The actual reference for the robot controller is obtained from the optimal sequence of kinematic states and controls $\{\mathbf{q}^*, \mathbf{u}^*\}$ in Section 3.4.1.

We opted to plan trajectories that are fully contained in a single 2D strain map, so that the output of our optimization would be intuitive for non-expert users and easy to visualize in two dimensions (e.g., on a computer screen). We selected the plane of elevation (PE) and shoulder elevation (SE) to be our controlled DoFs, while current values of axial rotation (AR) and muscle activation $\boldsymbol{a}$ are input for trajectory planning to select the correct strain map.

We define below the cost function and constraints of the OCP that allow us to set the optimality criterion with respect to which $\{\mathbf{q}^*, \mathbf{u}^*\}$ are found, and present how the OCP is solved in practice.

Cost Function

we design our cost function to capture important requirements for a safe rehabilitation trajectory: the movement should minimize the strain on selected tendons, produce low accelerations on the human body to reduce interaction forces and discomfort, and evolve towards the target final position. These elements are mirrored in the definition of our cost $L(\mathbf{x}, \mathbf{u})$, formalized as a weighted sum of the following terms:

- $L_\sigma = \sigma(\mathbf{x})$, accounting for the instantaneous (non-negative) strain that the selected tendon group undergoes;
- $L_{\text{acc}} = \dot{\mathbf{q}}^T \mathbf{Q_1} \dot{\mathbf{q}}$, accounting for the derivatives of our relevant state variables during the motion. The positive semidefinite matrix $\mathbf{Q_1} \in \mathbb{R}^{6\times6}$ selects only joint accelerations when planning on the current strain map;
- $L_\tau = \frac{1}{d_0^2}(\mathbf{q} - \mathbf{q}_\tau)^T \mathbf{Q_2} (\mathbf{q} - \mathbf{q}_\tau)$, accounting for the distance to the target human model pose $\mathbf{q}_\tau$, where d_0 represents the distance to the target pose at the beginning of the trajectory, to normalize the contribution of this term. Again, $\mathbf{Q_2} \in \mathbb{R}^{6\times6}$ is a positive semidefinite matrix selecting only the relevant human DoFs.

Overall, the cost to be minimized is:

$$L(\mathbf{x}, \mathbf{u}) = w_\sigma \sigma(\mathbf{x}) + w_{\text{acc}} \dot{\mathbf{q}}^T \mathbf{Q_1} \dot{\mathbf{q}} + w_\tau \frac{1}{d_0^2} \boldsymbol{\Delta}_\text{q}^T \mathbf{Q_2} \boldsymbol{\Delta}_\text{q} \tag{3.6}$$

where w_σ, w_{acc} and w_τ are scalars weighting the contribution of the various terms, and we indicated $(\mathbf{q} - \mathbf{q}_\tau) = \boldsymbol{\Delta}_\text{q}$ for brevity.

Constraints

We apply constraints to ensure that the OCP solution can be executed safely on the human subject. Their definition and role is detailed below.

- *Initial conditions*: every new trajectory $\{\mathbf{q}^*, \mathbf{u}^*\}$ should start from the current kinematic state of the human model $\hat{\mathbf{q}}_{\text{curr}}$, estimated through the robot's encoders (see Section 3.3.2). This value is used as the initial condition for the next instance of the OCP enforcing $\mathbf{q}(t = 0) = \hat{\mathbf{q}}_{\text{curr}}$, where the initial time is set to be 0 without loss of generality.
- *Torque limits*: to ensure that following the optimal trajectory does not require excessive torque to be exerted along the human DoF, we limit them within heuristic bounds by requiring $\mathbf{u}_{\text{min}} \leq \mathbf{u}(t) \leq \mathbf{u}_{\text{max}}$, where $\mathbf{u}_{\text{min}}, \mathbf{u}_{\text{max}} \in \mathbb{R}^3$. Bounds can be adjusted to enforce different torque limits on the various DoF accounting, for example, for the fact that torques along SE should counteract gravity.
- *Terminal condition*: to impose an acceptable final state for the human model, we prescribe that $(\mathbf{q}_{t=T_f} - \mathbf{q}_\tau) \odot (\mathbf{q}_{t=T_f} - \mathbf{q}_\tau) \leq \boldsymbol{\epsilon}$, where T_f represents the length of the planning horizon, and $\odot$ the Hadamard product. If reaching the target final state is possible in the current OCP instance (e.g. if T_f is long enough), the

parameter $\boldsymbol{\epsilon} \in \mathbb{R}^6$ is defined for every element of $\mathbf{q}$, meaning both final pose and velocity of the human model can be specified fully. Otherwise, if the target pose specified by $\mathbf{q}_T$ is too far from $\hat{\mathbf{q}}_{\text{curr}}$, only the final velocities are limited. This choice guarantees low velocities of the human arm at the end of the planning horizon and ensures recursive feasibility through a safe terminal set $\mathbb{S}_T$ when the optimization is performed iteratively over a receding time horizon.

Optimal control problem transcription

Overall, the optimal control problem that we aim to solve reads as follows:

$$\begin{aligned} &\min_{\mathbf{q}(\cdot),\mathbf{u}(\cdot)} \int_0^{T_f} L\big(\mathbf{x}(t),\mathbf{u}(t)\big)\,dt \\ &\text{subject to:} \\ &\quad \mathbf{q}(0) - \hat{\mathbf{q}}_{\text{curr}} = \mathbf{0} && \text{initial state} \\ &\quad \mathbf{u}_{\min} \le \mathbf{u}(t) \le \mathbf{u}_{\max} && \text{torque limits} \\ &\quad \mathbf{q}_{t=T_f} \in \mathbb{S}_T && \text{terminal set} \end{aligned} \tag{3.7}$$

To solve this OCP, we cast it into discrete time, transcribing equation (3.7) into an equivalent Non-Linear Programming problem (NLP) that can be solved by structure-exploiting solvers. We used orthogonal collocation techniques [150], hence approximating the state trajectories with suitable d^{th}-order polynomial splines. The overall optimization horizon T_f is broken down into N intervals of equal length, and inside the generic interval $[t_k, t_{k+1}]$ we select $d = 3$ Legendre-Gauss collocation points at which the dynamics as in equation (3.5) are enforced. This step results in additional continuity and collocation constraints that need to be considered in the NLP formulation. Resulting optimal trajectories $\{\mathbf{q}^*, \mathbf{u}^*\}$ were upsampled for smoother execution.
We implemented the resulting NLP in Python and solved it with IPOPT [151] leveraging MA27 [152] as a linear solver, while derivatives were provided by CasADi [149] and OpenSimAD. To clarify the computational advantage of using the differentiable skeletal dynamics, we also embedded numerically evaluated gradients with perturbation calls to the original OpenSim model directly in the collocation constraints of the NLP, through CasADi. We ran simulations for both cases to compare computational requirements.

3.3. Dynamic Adaptation to Human Actions

The trajectory on strain maps generated by the planner (Section 3.2.2) provides a reference for the robot controller. Nevertheless, unpredictable volitional actions of the patient and/or reflexive reactions to physical human-robot interaction can induce large interaction forces and corresponding muscle activations that lead to potentially large changes in the strain map. In this section, we describe how we exploit the robot's sensory system to make real-time updates on estimates of patient pose and muscle activity. These estimates are used to update the strain maps, which the planner will use to re-plan the robot trajectories in an online manner to avoid dangerous high-strain areas and abide by given optimization objectives.

3.3.1. Online estimation of muscle activity

Estimating the activity of deep rotator cuff muscles is not feasible using traditional surface electromyography (EMG) measurements. To overcome this limitation, we employed a model-based estimation approach to infer individual muscle activations within our biomechanical model during robot-assisted rehabilitation. Specifically, we used the Rapid Muscle Redundancy (RMR) solver [59, 147], which yields physiologically plausible muscle activations in over-actuated systems — i.e., systems with more muscles than degrees of freedom — while accounting for the stabilizing role of the rotator cuff at the glenohumeral joint [59]. As the human musculoskeletal model is inherently redundant, the solver leverages nonlinear optimization to estimate muscle activations consistent with observed kinematics and forces. The RMR solver minimizes a muscle effort cost function subject to physiological constraints to generate realistic recruitment patterns.

The original open-source implementation of RMR was designed for offline use. In contrast, our application requires real-time estimation, with the model state updated continuously based on measured joint kinematics (Section 3.3.2) and external human-robot interaction wrenches (Section 3.4.2). To this end, we restructured the solver into a class-based module integrated with the Robot Operating System (ROS), enabling online estimation of muscle activations $\hat{\boldsymbol{a}}$ in a dedicated sub-process. These real-time estimates were used to select the tendon strain map that informed the trajectory optimization described in Section 3.2.2.

The code implementation is available in the project repository at https://github.com/itbellix/baton-robotic-rehab.

3.3.2. Human kinematics estimation

We estimated the human model's kinematic state $\hat{\mathbf{q}}$ through the position, orientation, and twist of the robot's end-effector (EE), which we mapped to the shoulder state $\mathbf{x}$ according to the fixed rigid connection between the EE and elbow locked at 90°(see Figure 3.3). An estimation of the full human body pose is out of the scope of our current work, so we assume for simplicity that the center of the glenohumeral joint is stationary during the experiments and that the orientation of the human torso is fixed. Under these conditions, monitored during the experiments, it is possible to reconstruct the shoulder kinematic state through the robot's encoders.

We observe that the orientation of the human elbow expressed in the reference frame of the human shoulder is[1] :

$$^{sh}\mathbf{R}_{elb} = {}^{sh}\mathbf{R}_{base}\,{}^{base}\hat{\mathbf{R}}_{EE}(t)\,{}^{EE}\mathbf{R}_{elb}, \tag{3.8}$$

where *sh* denotes the shoulder frame, *base* the frame attached to the robot's base, and *elb* the frame fixed at the human elbow (see Figure 3.3). The first and last rotations are known and fixed, while ${}^{base}\hat{\mathbf{R}}_{EE}(t)$ can be estimated through the robot's forward

[1] we denote as ${}^{I}\mathbf{p}_k = [{}^{I}p_k^x, {}^{I}p_k^y, {}^{I}p_k^z]^T \in \mathbb{R}^3$ the position of point k w.r.t. frame "I", and ${}^{J}\mathbf{R}_I \in SO(3)$ the rotation matrix to express it in frame "J".

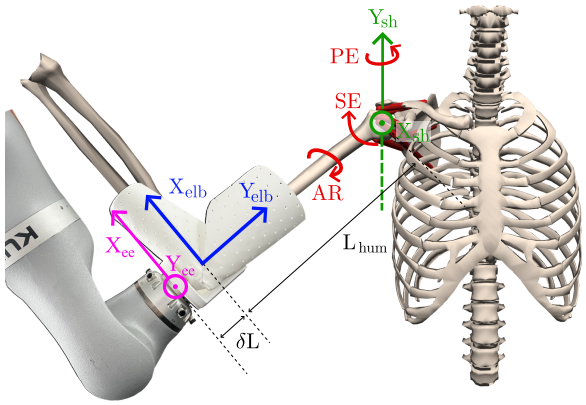

Figure 3.3: *The human and robot coordinate systems and their correspondence. The shoulder frame (green) originates in the center of the glenohumeral joint. The glenohumeral joint DoFs (PE, SE, and AR) are shown in red. The elbow frame (blue) originates in the center of the elbow, and a fixed transformation relates it to the robot's end-effector frame (pink).*

kinematics. Moreover, following the definition of the glenohumeral DoFs in the human model [58], we can write:

$$^{sh}\mathbf{R}_{elb} = \mathbf{R}_Y(PE)\mathbf{R}_X(-SE)\mathbf{R}_Y(AR). \tag{3.9}$$

We first obtain our estimates $\hat{PE}$ and $\hat{SE}$ leveraging the end-effector position $^{base}\hat{\mathbf{p}}_{EE}(t)$, and then estimate $\hat{AR}$ through equation (3.9). Formally, we define a unit vector $\boldsymbol{\beta} \in \mathbb{R}^3$ expressing the direction of the humerus (upper arm) pointing from the center of the glenohumeral joint to the human elbow, and use its Cartesian components in the shoulder frame to estimate the human state:

$$\begin{aligned} \boldsymbol{\beta} &= {}^{sh}\mathbf{R}_{base}\frac{{}^{base}\hat{\mathbf{p}}_{EE}(t) - {}^{base}\mathbf{p}_{GH}}{||{}^{base}\hat{\mathbf{p}}_{EE}(t) - {}^{base}\mathbf{p}_{GH}||_2^2} \\ \hat{PE} &= \operatorname{atan2}(\beta^x, \beta^z) \\ \hat{SE} &= \arccos\left(\boldsymbol{\beta} \cdot [0 \;\; -1 \;\; 0]^T\right) \end{aligned} \tag{3.10}$$

where $^{base}\mathbf{p}_{GH}$ is the known position of the center of the glenohumeral joint. Velocities of the generalized DoFs in the human model are estimated through the body twist of the last link of the robot, obtained through the robot's Jacobian:

$$^{base}\hat{\boldsymbol{\nu}}_{EE} = [^{base}\hat{\mathbf{v}}_{EE}^T \quad ^{base}\hat{\boldsymbol{\omega}}_{EE}^T]^T \in \mathbb{R}^6. \tag{3.11}$$

Since the robot's EE and the human elbow are rigidly attached to each other, we can transform $^{base}\hat{\boldsymbol{\nu}}_{EE}$ into the twist $^{sh}\hat{\boldsymbol{\nu}}_{elb}$ through the known transformation:

$$^{sh}\hat{\boldsymbol{\nu}}_{elb} = \begin{bmatrix} \mathbf{R} & 0 \\ 0 & \mathbf{R} \end{bmatrix} {}^{base}\hat{\boldsymbol{\nu}}_{EE} \tag{3.12}$$

where $\mathbf{R} = {}^{\text{sh}}\mathbf{R}_{\text{elb}}\,{}^{\text{elb}}\mathbf{R}_{\text{EE}}\,{}^{\text{EE}}\mathbf{R}_{\text{base}}$. Having retrieved $\hat{PE}$ and $\hat{SE}$, ${}^{\text{sh}}\hat{\boldsymbol{v}}_{\text{elb}}$ can be further transformed into ${}^{\text{sh'}}\boldsymbol{v}_{\text{elb}}$, denoting the elbow twist expressed in the intermediate rotated frame where *SE* is defined (which we name *sh'*), and into ${}^{\text{sh''}}\boldsymbol{v}_{\text{elb}}$, which denotes the elbow twist expressed in the frame where *AR* is defined (which we name *sh''*). Overall, we can obtain the velocity estimates for the human shoulder model as follows:

$$\begin{aligned}
\dot{\hat{PE}} &= {}^{\text{sh}}\omega^{y}_{\text{elb}}, \quad \text{with} \quad {}^{\text{sh}}\boldsymbol{\omega}_{\text{elb}} = \frac{{}^{\text{sh}}\mathbf{p}_{\text{elb}} \times {}^{\text{sh}}\hat{\mathbf{v}}_{\text{elb}}}{L^2_{\text{tot}}} \\
\dot{\hat{SE}} &= -{}^{\text{sh'}}\omega^{x}_{\text{elb}}, \quad \text{with} \quad {}^{\text{sh'}}\boldsymbol{\omega}_{\text{elb}} = \frac{{}^{\text{sh'}}\mathbf{p}_{\text{elb}} \times {}^{\text{sh'}}\hat{\mathbf{v}}_{\text{elb}}}{L^2_{\text{tot}}} \\
\dot{\hat{AR}} &= {}^{\text{sh''}}\hat{\omega}^{y}_{\text{elb}}
\end{aligned} \tag{3.13}$$

where $L_{\text{tot}} = L_{\text{hum}} + \delta L$ is the distance between the center of the glenohumeral joint and the robot's EE (see Figure 3.3), ${}^{\text{sh}}\mathbf{p}_{\text{elb}}$ is the position of the human elbow expressed in the shoulder frame, and ${}^{\text{sh'}}\mathbf{p}_{\text{elb}}$ is the equivalent quantity expressed in the *sh'* frame.

Finally, we differentiated numerically $\hat{PE}$, $\hat{SE}$, and $\hat{AR}$ to obtain the corresponding acceleration vector $\ddot{\boldsymbol{\theta}}_{\text{S}}$. Exponential filtering was applied to the estimated values to reject high-frequency components incompatible with human movements in physiotherapy.

3.4. Robot control and experimental design

To test BATON as a novel approach to robotic-assisted rotator cuff rehabilitation and demonstrate its capabilities, we designed and performed various lab-based experiments. These proof-of-concept scenarios are essential for the validation of our approach as a whole, testing BATON against the state of the art and in situations where its dynamic adaptation becomes crucial for guaranteeing safe rehabilitation.

3.4.1. Robot control

In addition to human state estimation, our collaborative robotic arm administers therapeutic motion to the human subject. We map the optimal trajectory $\mathbf{q}^*$ resulting from equation (3.7) into the corresponding robot's end-effector reference pose:

$${}^{\text{base}}\bar{\mathbf{R}}_{\text{EE}} = {}^{\text{base}}\mathbf{R}_{\text{sh}}\,{}^{\text{sh}}\mathbf{R}_{\text{elb}}(\mathbf{q}^*)\,{}^{\text{elb}}\mathbf{R}_{\text{EE}} \tag{3.14}$$

$${}^{\text{base}}\bar{\mathbf{p}}_{\text{EE}} = {}^{\text{base}}\mathbf{p}_{\text{GH}} + {}^{\text{base}}\bar{\mathbf{R}}_{\text{EE}}\,{}^{\text{EE}}\mathbf{p}_{\text{GH}} \tag{3.15}$$

Equations equation (3.14)-equation (3.15) ensure that the robot's end-effector pose is consistent with the desired trajectory for the human. Given the optimal shoulder angles, the EE pose follows from the human elbow position and orientation, and is oriented such that the upper arm always points to the center of the human shoulder.
To track target end-effector poses, we used an impedance controller that governs the interaction force/torque in Cartesian space as:

$$\mathbf{f} = \mathbf{K}(\bar{\mathbf{x}}_{\text{EE}} - \hat{\mathbf{x}}_{\text{EE}}) - \mathbf{D}\dot{\hat{\mathbf{x}}}_{\text{EE}}, \tag{3.16}$$

where $\mathbf{f}$ is the interaction force/torque between the robot and its patient, $\bar{\mathbf{x}}_{EE} \in \mathbb{R}^6$ is the commanded EE reference retrieved from ${}^{base}\bar{\mathbf{R}}_{EE}$ and ${}^{base}\bar{\mathbf{p}}_{EE}$, and $\hat{\mathbf{x}}_{EE} \in \mathbb{R}^6$ is the measured actual pose of the robot, respectively. $\mathbf{K} \in \mathbb{R}^{6\times6}$ and $\mathbf{D} \in \mathbb{R}^{6\times6}$ are the Cartesian stiffness and damping matrix, respectively. The impedance controller operates at 200 Hz.

In addition to the Kuka's built-in gravity compensation, we include further compensation to support the mass of the subject's arm. Specifically, we use the optimal generalized torques from the OCP in equation (3.7) to adjust the vertical reference for the robot's EE, increasing it proportionally to the torques required to produce the optimal movement. This adjustment, further detailed in the Appendix, allows our system to keep the actual vertical position of the human arm close to the prescribed one, leveraging pure impedance control without the need to feed-forward assistive torques.

3.4.2. Experimental design for evaluation of BATON

Our experimental setup consisted of a KUKA LBR iiwa 7 robotic arm instrumented with a Bota SensONE force-torque (F/T) sensor (Bota Systems AG, Zurich, CH) and a custom thermoplastic elbow brace for interfacing the human subject with the robot. The brace was molded to cradle the subject's elbow (Figure 3) and mounted directly onto the F/T sensor at the robot's end-effector, fixing the transformation between the robot EE frame and the human elbow frame. This configuration allowed the robot to impose motion on the human arm, support its weight, and simultaneously estimate the subject's biomechanical state.

A healthy adult subject participated in our lab-based experiments, approved by the Human Research Ethics Committee at Delft University of Technology. The brace was strapped to the participant's arm, and they were instructed to maintain a fixed torso posture relative to the robot base, satisfying the assumption of a stationary glenohumeral joint center required by our state estimation model.

To evaluate BATON and highlight the importance of modeling human arm dynamics in trajectory planning, we examined two rehabilitation scenarios: (1) the early stages of therapy, during which the robot provides full guidance while the subject remains passive; and (2) the later stages, where the subject progressively regains strength and actively participates in the motion. The first scenario also served to benchmark BATON against the state-of-the-art (SoTA) method proposed by Prendergast et al. [144], which generates strain-aware shoulder rehabilitation movements based on tendon behavior. Their method uses a modified A* algorithm to traverse discrete passive strain maps while minimizing cumulative strain, returning a kinematic sequence of optimal joint angles $\mathbf{q}^*_{SoTA}$. We implemented this planner in our framework and compared its performance with BATON's in guiding a passive subject. The second scenario, involving an active subject, falls outside the applicability of the SoTA method, which assumes constant muscle activations and thus cannot accommodate unpredictable activation changes.

Passive human subject

In this scenario, the subject remained relaxed, allowing the robot to lead the movement. Muscle activity was negligible, and tendon strain was primarily determined by the

shoulder pose. To ensure uniformly low strain across the rotator cuff, we monitored a strain map that aggregated the maximum strain values over the supraspinatus, infraspinatus, subscapularis, and teres minor tendons. The strain value at a given model state $\mathbf{x}$ was defined as the maximum of these individual tendon strains. To maintain consistency on the strain map, the axial rotation (AR) degree of freedom (DoF) was locked, so that the complete movement toward the target arm state $\mathbf{q}_T$ evolved along a single map. Negligible muscle activation and constant axial rotation imply that trajectories can be optimized and then executed in an open loop. Multiple target states $\mathbf{q}_T$ were defined and used as inputs to BATON to explore a broader range of motion at the glenohumeral joint. Each optimization instance used a time horizon of T_f = 5 seconds, discretized into N = 50 intervals. Cost function weights were selected based on extensive simulations to provide parameter tuning guidelines.

All candidate trajectories were first tested in the Gazebo simulator to verify the robot's ability to execute them accurately. We then compared BATON-generated trajectories to those from the SoTA method, each executed on the physical robot. The SoTA trajectory $\mathbf{q}^*_{\text{SoTA}}$ was upsampled to match the execution frequency of BATON's output and fed to the controller via equations (3.14) and (3.15). Performance was evaluated in terms of strain minimization, trajectory smoothness, and human–robot contact forces. Both methods started from the same human shoulder state $\mathbf{q}(0)$ and navigated between the same targets.

Active human subject

Due to human volitional actions and/or reflexive reactions during therapy, the level of muscle activity and forces can change quickly and significantly, impacting tendon strain estimates. BATON needed to adapt its trajectory planning in real time to maintain safety under these dynamic conditions. We shortened the planning horizon to T_f = 1 second, with N = 10 intervals, allowing for continuous replanning using a receding horizon scheme.

We first used simulations to investigate the reactiveness of BATON to artificial muscle activation changes when treating a specific tendon. Without loss of generality, the infraspinatus was considered. The initial human pose was set at $SE_0 = PE_0 = 60^\circ$, while the target position to $\bar{PE} = 45^\circ$, $\bar{SE} = 95^\circ$, requiring movement along multiple human DoFs. Simulation also allowed us to consider drastic variations on fabricated maps to trigger more extreme trajectory adaptations as a consequence, for example, of a strong reflex response. In this case, the previous initial and target positions were inverted, representing the previous movement that evolved in the opposite direction.

Then, we moved to a closed-loop execution of BATON on the real robot, employing realistic infraspinatus activity estimation $\hat{a}_{\text{is}}$ through the RMR solver (Section 3.3.1).

In this scenario, we analyzed the effect that physical human-robot interaction and volitional human behavior have on online therapy adjustments. BATON guided our subject between the same initial and final poses ($SE_0 = PE_0 = 60^\circ$, $\bar{PE} = 45^\circ$, $\bar{SE} = 95^\circ$) under two conditions: (i) the subject complied with the robot, and (ii) the subject actively engaged the infraspinatus by actively externally rotating their arm. This rotation, while not substantially disrupting the robot-guided trajectory, was resisted by the Cartesian controller, which maintained the initial value of the *AR* DoF. The resulting interaction force

elicited infraspinatus activation, consistent with its primary anatomical function [153]. Model-based computations were run on a Dell Latitude 7420 laptop with an i7-1185G7 CPU, interfaced with a workstation running a Xeon W-2123 and executing the robot's impedance controller. This ensured low-level control at 200 Hz, decoupled from potential delays in high-level planning. The Robot Operating System managed communications between the various software modules.

3.5. Results

We performed a thorough evaluation of the proposed methods. Regarding the planning aim (A1), we analyzed the effect of optimization parameters on the planned trajectory (Section 3.5.1) and compared BATON to the SoTA method in terms of the quality of the executed paths (Section 3.5.2). Regarding adapting to changing muscle activity (A2), we demonstrated BATON's ability to re-plan online based on changing strain maps, first in simulation and then thanks to real-time muscle activity estimates (Section 3.5.3). We concluded by conducting an analysis of real-time computational performance (Section 3.5.4).

3.5.1. Effect of OCP parameters on trajectories (passive subject)

We first ran BATON in simulation to examine the effects of optimization parameters and to determine the best relative cost function weights for synthesizing rehabilitation movements for passively assisted patients. We explored various permutations of the cost function weights w_σ, w_{acc}, and w_τ, modulating the importance that tendon strain, human accelerations, and movement towards the goal pose have in the overall cost function. Figure 3.4 shows some of the representative scenarios. We fixed $w_{\text{acc}} = 10$ and varied the other two parameters to achieve rather extreme cases. If minimizing the strain is ignored ($w_\sigma = 0$), the planner would choose the shortest path (dashed blue line). As higher weights w_σ are used, the trajectory makes a wider excursion in the low-strain regions (solid and dotted blue lines). While this guarantees safer therapy, we observed that increasing this term gradually produces wider trajectories while the quantitative decrease in strain becomes negligible. In other words, wider excursions do minimize the cost function value, but they lead to insignificant reductions in strain. Given this trade-off, we fixed $w_\sigma = 1$ for later experiments.

With a passive subject, the strain map was fixed and did not vary with time. Thus, we could directly optimize the entire trajectory and execute it in an open-loop manner. Here, both final human pose and velocities are constrained (see Section 3.2.2), so the target pose is reached even when $w_\tau = 0$. This terminal constraint effectively simplifies the tuning of our cost function to only two parameters in this case. By selecting $w_{\text{acc}} = 10$, $w_\sigma = 1$ the resulting trajectory traverses safe low-strain regions ($\sigma \leq 2\%$). These weights were used for the robot experiments in Section 3.5.2 below.

However, enforcing a terminal constraint on the final human pose is undesirable with shorter planning horizons, which are needed to react to changes in the strain map as muscle activations vary (active subject scenario, Section 3.5.3). A trajectory deemed

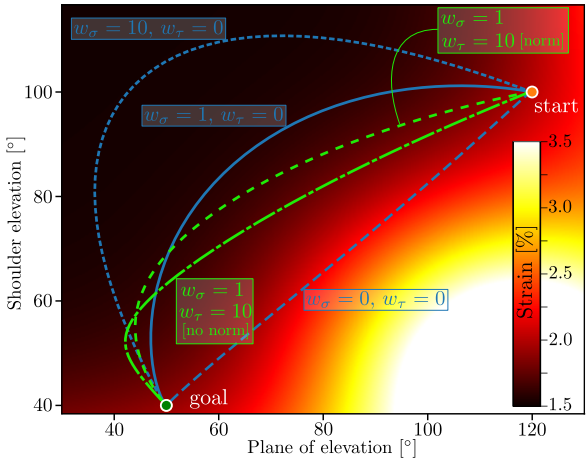

Figure 3.4: *Effect of tuning the cost function parameters on the optimal trajectories for the human DoF. In blue, the influence of w_σ on the optimal path. In green, the effect of normalizing L_τ when weighting the distance to the target final pose. To guarantee slow motions, we kept w_{acc} = 10.*

safe on the initial map may later intersect high-strain regions as the map is updated. This motivates setting $w_\tau \neq 0$, and investigating the effect of normalizing the distance to the target L_τ on the optimal trajectory (Figure 3.4, green lines). When L_τ was normalized by the initial distance from $\mathbf{q}_\tau$, lower-strain trajectories were produced, prioritizing strain minimization over approaching the goal earlier. This normalization ensures that the goal pose provides a movement direction, and prevents this information from vanishing when the goal is closer or dominating the cost when $\mathbf{q}_\tau$ is farther.

3.5.2. Quality of executed path w.r.t. SoTA

Resulting trajectories from BATON were compared to the ones obtained with the SoTA planner from the literature [144]. Figure 3.5 shows the experimental setup and the target poses on the strain map, together with the optimal trajectories that were delivered to our test subject by the robot in both cases. For both executions, the robot started at the shoulder pose marked by the blue dot in the low-strain zone. First, it moved the patient to the pose marked by the orange dot and then toward the state indicated by the green dot to increase the level of strain in a controlled manner and augment the subject's range of motion. Finally, the robot moved the subject back into the low-strain zone denoted by the pink dot.

BATON's resulting trajectory avoided higher strain regions (Figure 3.5, top left), navigating the rather large range of motion with lower strain overall. On the other hand, the SoTA method produced trajectories that tend to cross zones of higher strain, and lead to oscillations in the subject's movement (especially evident towards the last pose). The bottom row of Figure 3.5 (left) shows the smoothness of the resulting movements in terms of human acceleration magnitude, with the SOTA approach producing consistently

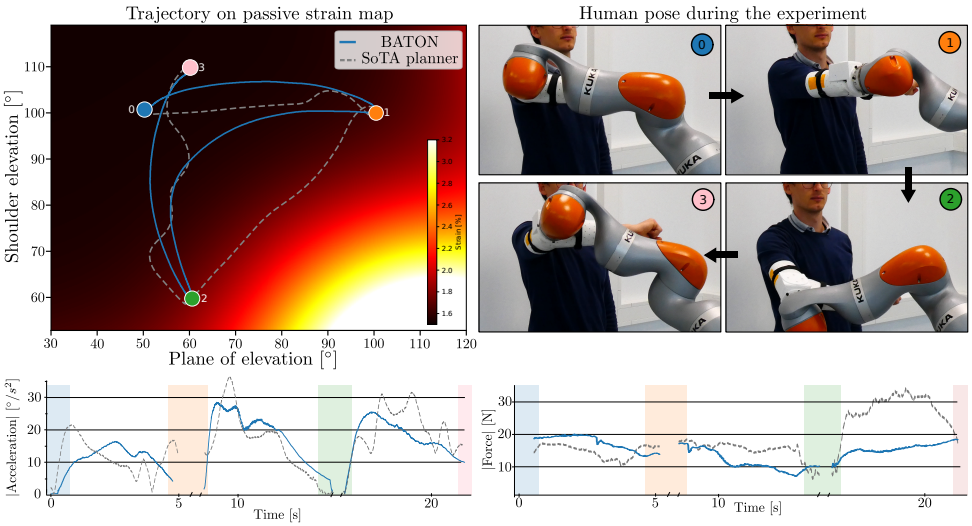

Figure 3.5: *Top left: human trajectories produced by our approach for passive rehabilitation (solid blue), visualized on the corresponding strain map together with the result of the planner in [144] (dashed grey). The two lines are the actual trajectory, in human coordinates, along which the robot moved our subject. The paths connect target human poses given as goals to 3 instances of our trajectory optimization problem minimizing the maximum strain that the rotator cuff tendons encounter, and 3 instances of the planner in [144]. Top right: goal poses of the test subject during the experiment. The graphs at the bottom display the magnitude of the acceleration along the two trajectories (left) and the interaction force magnitude between the robot and subject during the motion (right). Note that the data corresponding to re-planning periods is not shown for conciseness.*

higher accelerations with respect to BATON. The evolution of the interaction forces between our subject and the robot is also shown for both approaches (bottom right), with BATON's trajectory requiring lower interaction forces, especially during the last part of the movement.

3.5.3. Online adaptations (active subject)

Simulations

First, we investigated the responsiveness and adaptability of our control scheme when the strain maps change as a consequence of artificial changes in infraspinatus activations. Owing to the shorter time horizon over which we ran our trajectory optimization, BATON was able to achieve a trajectory re-planning with an update rate of roughly 10 Hz. This permits us to generate new trajectories that can account for strain map variations, allowing us to control human movement in a closed-loop manner.
The effect of two different muscle activation changes is visible in Figure 3.6: a sudden step increase in activation (white trajectory) and a gradual decrease in activation (green

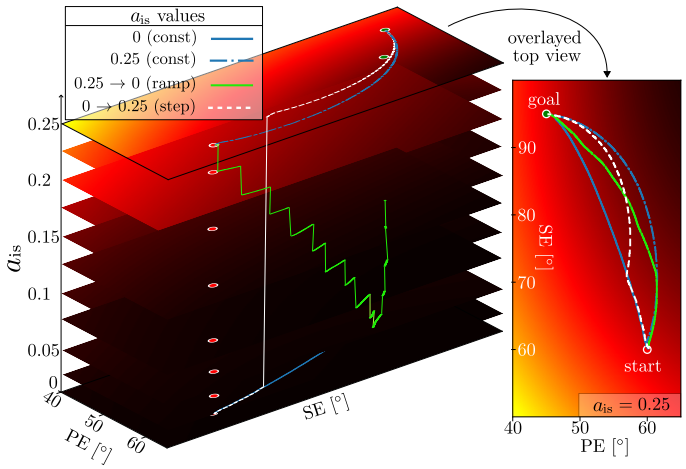

Figure 3.6: *The simulated effects of infraspinatus muscle activation (a_{is}) on the optimal trajectory with the same target pose, when a_{is} follows different activation profiles (green: decreasing ramp, white: step) or remains constant (blue). Large activation steps are used for ease of visualization, and all of the resulting trajectories are projected on the strain map corresponding to the highest a_{is} for comparison to the right. Our musculoskeletal trajectory optimization is capable of accounting for sudden variations in the strain level, as a consequence, e.g., of varying muscle activation.*

trajectory). These two cases are shown alongside the two boundary cases where activation is kept constantly high or low (resulting in the blue trajectories in the figure). We visualize all of the trajectories jointly on a single strain map, representing the strain map defined by the highest activation, set at a_{is} = 0.25 here (Figure 3.6, right). Re-projecting all of the trajectories on the same plane highlights differences due to the gradual and sudden variations of the activation level, bounded by the two cases in which the activation was kept constant.

To understand the effects of more sudden and large changes in activation, for example, due to a strong reflex response, we tested cases in which BATON is presented with larger, more localized, and sudden variations in the strain maps (see Figure 3.7). First, the strain map on which to plan has low strains, so the optimal trajectory points directly to the goal (Figure 3.7, left). We then introduced a sudden, large local variation in the strain map topography, so that higher strain appears along the direction of movement. BATON reacted to the variation by re-planning a new trajectory that keeps evolving towards the lower-strain region, deviating from the older plan proposed at the previous time step (white dots in the figure). Similarly, we introduced a second sudden large local variation in the strain map, observing a new trajectory adaptation that allows the simulated movement to remain safe.

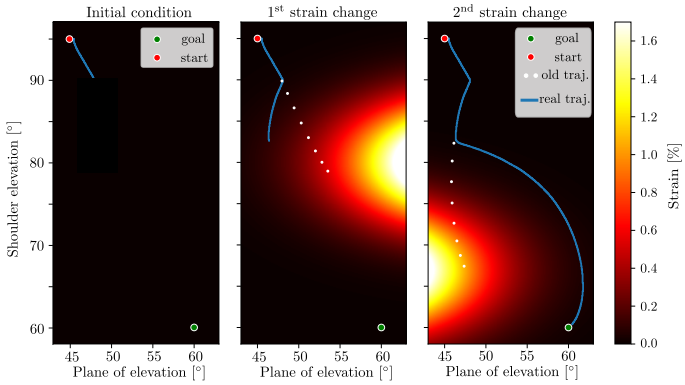

Figure 3.7: *Real-time trajectory optimization allows to account for sudden variations in the strain level. We show here how the optimal trajectory changes as a consequence of rapid variations of (fabricated) strain maps. The trajectory optimization starts on a completely safe strain map, generating a motion directly toward the goal (left). Sudden variations of the maps cause the reference to avoid higher-strain zones (center and right), generating a trajectory that deviates from the shortest path adapting to the new scenarios, and prioritizing safety.*

Real system

Having explored and compared the real-time adaptation of the reference optimal rehabilitation trajectories in simulation, we proceeded to demonstrate this in a real robotic-assisted experiment. Here, online estimates for all of the muscles in the human model were achieved through the RMR solver (Section 3.3.1) at 20 Hz. Based on the estimated infraspinatus activation $\hat{a}_{is}$, the strain map considered by trajectory optimization was dynamically updated at runtime. The subject followed the robot-led movement twice, between the same starting and goal shoulder poses. In the first condition, they remained passive as the robot moved their arm, and the strain map topography was mostly static during execution (Figure 3.8, first row). In the second condition instead, they were instructed to perform an external rotation of their arm, resisted by the robot, to elicit infraspinatus activation (Figure 3.8, second row). Selected snapshots of the two movements, visualized on the time-varying strain maps, reveal BATON's adaptations to muscle activation changes. In particular, the two executions lead to a difference of about 10° in human poses, as a result of adaptations of the human movements to traverse lower tendon strain areas defined by the time-dependent muscle activity values (Figure 3.8, third row) resulting from increased human-robot contact wrenches (Figure 3.8, bottom row).

3.5.4. Computational performances

Finally, we used simulations to compare the computational efficiency of our choice to capture the human skeletal dynamics with OpenSimAD, against numerically querying the original OpenSim model during the optimization (Section 3.2.1). This was done

for an active subject as well as for a passive subject, modulating the time horizon for the trajectory optimization instances equation (3.7). The average computation times for solving equation (3.7) when the system dynamics are enforced at the collocation points through the two different methods indicate a difference of roughly two orders of magnitude in favor of BATON (Tab. 3.1). Real-time trajectory optimization was possible only with OpenSimAD, granting re-planning capabilities at a frequency of roughly 10 Hz with appropriate selection of the planning horizon and its discretization. BATON achieved the same computational performances during the robot experiments.

Table 3.1: *Computation time to solve a single OCP instance with different time horizons: A) passive subject, B) active subject*

Scenario	OpenSim+CasADi	OpenSimAD (BATON)
A (T_f = 5 s, N = 50)	227.0 s	1.9 s
B (T_f = 1 s, N = 10)	8.50 s	0.12 s

3.6. Discussion

3.6.1. Motion Planning on strain maps

BATON is a new approach that enables robot-assisted physiotherapy to both automate and personalize rehabilitation by navigating underlying human biomechanical outcomes during therapy delivery. Achieving our first research aim (A1), BATON directly embeds human biomechanics into a robot controller through trajectory optimization and adapts to muscle activation changes as they occur. Accounting for high-fidelity musculoskeletal outcomes is crucial to deliver effective robotic rehabilitation. We developed and applied BATON to address the complexities and challenges of traditional physical therapy for rotator cuff injuries, where insights into patient-level tendon strains are paramount. Knowledge of the patient's tissue loads and range of motion is essential for optimizing safe movements that avoid high-strain configurations associated with risks of re-injury. These biomechanical insights allow us to plan and execute robotic rehabilitation trajectories, but can also be harnessed by physical therapists to guide the efficacy and safety of applied movements.

In both passive and active scenarios, we found BATON to operate several orders of magnitude faster than using traditional musculoskeletal modeling and simulation techniques (Tab. 3.1). As a result, real-time re-planning of reference trajectories for the rehabilitation robot was attained at an average update rate of just below 10 Hz.

3.6.2. Dynamic adaptation to the human actions

An integral part of the safety and potential autonomy of BATON is the feedback of the human response into the planning and control loop. We demonstrate the estimation and adaptation to changes in muscle activity in both simulation (Figs 3.6-3.7) and activity inducing movements in a test subject (Figure 3.8). The online model-based estimates of muscle activity account for changes in the patient's unpredictable volitional

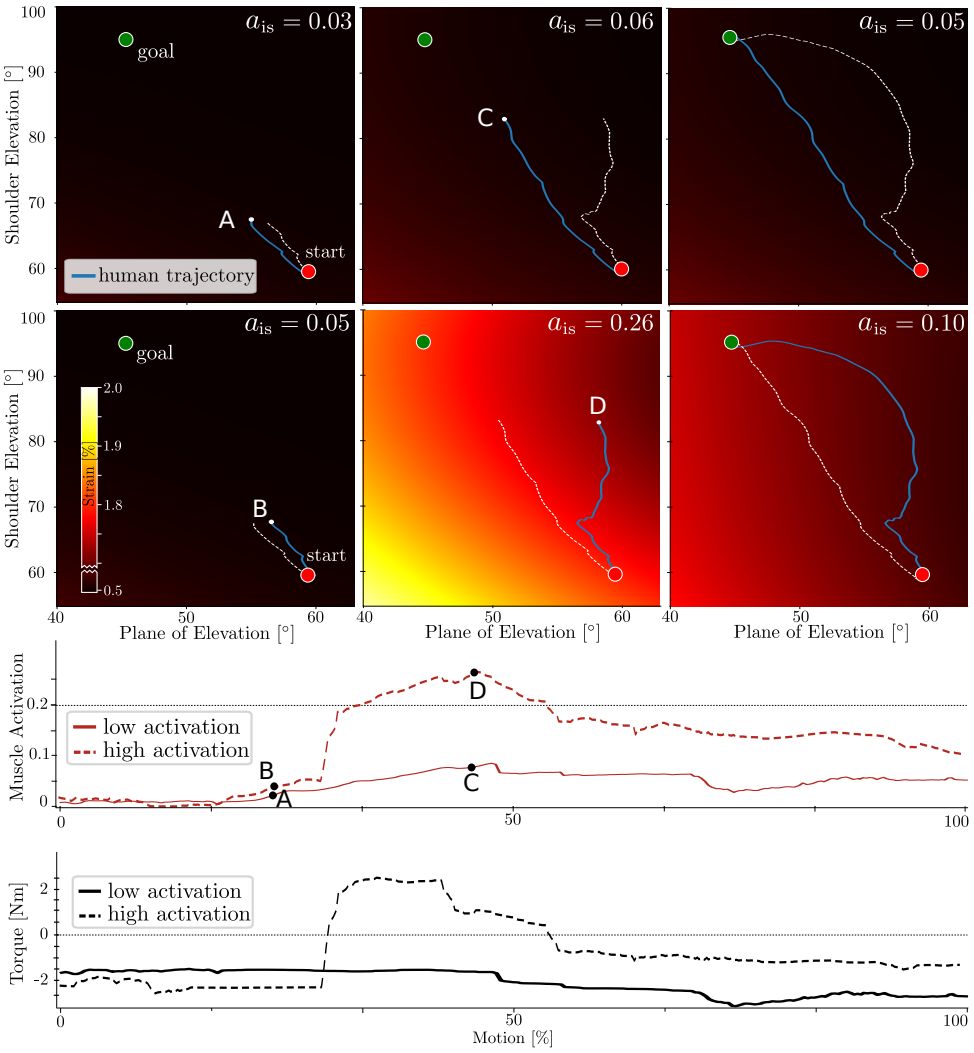

Figure 3.8: *Different infraspinatus activation a_{is} causes different strain map topography for the corresponding tendon in our physical human-robot interaction experiments. Feeding human pose and interaction torque to BATON, it adapts to changes in estimated muscle activity. Here, we reveal the progression of rehabilitation trajectories along varying strain maps for our subject complying with the robot guidance (top row) and voluntarily activating their infraspinatus (second row). The solid blue lines represent the trajectory executed for each case, while dashed white lines allow us to compare them to the other case. Bottom rows: evolution of estimated infraspinatus activation $\hat{a}_{is}$, and interaction torque modulated by the subject during the motion for both cases.*

actions and/or reflexive responses due to physical human-robot interaction, achieving our second aim (A2). As a consequence, our system can inform the planner about underlying muscle activity, e.g., rotator cuff muscles, which cannot be measured by surface EMG. Although the proposed approach does not require surface EMG, such data could be used, if available, to augment the solver in estimating muscle activity.
We analyzed our algorithm over multiple edge cases in simulation (Figure 3.6 and Figure 3.7), to verify the capability of our online trajectory optimization to cope with extreme strain map variations (potentially due to muscle activation changes). Then, we leveraged a model-based muscle estimation algorithm to close the loop between sensing, online trajectory optimization, and reference tracking on the real robot (Figure 3.8). The use of the RMR solver enabled our system to estimate physiological muscle recruitment without the use of instrumenting our subject with any kind of sensor (e.g., surface EMG ones), decreasing the time required for setup and enabling simultaneous monitoring of virtually any muscle in the human body (even deep stabilizers like the rotator cuff). Considering the infraspinatus muscle as a target tissue for our real-robot experiments, we achieved online therapy adaptations to its activation changes during physical human-robot interaction (Section 3.5.3). Muscle activation was monitored both when the subject complied with the robot motion commands, and when they voluntarily elicited muscle activation. In the first case, rather small variations occurred, and they were likely due to the mechanical role of the infraspinatus as a glenohumeral stabilizer during humeral elevation. In the second case, voluntary muscle activation was estimated, in a scenario similar to real therapy when a patient's neuromuscular reflexes might be triggered, or the human might interact in unexpected ways with the robot.

3.6.3. Evaluation insights of BATON performance

Further, we analyzed the performances of BATON through simulations and physical experiments and achieved the third aim of this study (A3). The optimized trajectories can be tuned to reach a trade-off between the minimization of tendon strain and the directness of the path (Figure 3.4). BATON's optimized trajectories guaranteed lower instantaneous and cumulative tendon strain than the SoTA planner [144], which instead traversed higher strain regions (Figure 3.5). Moreover, BATON produced smoother trajectories, while the SoTA planner delivered trajectories presenting higher oscillations, potentially in contrast with natural human movement preferences [154]. These oscillations are intrinsic to the assumptions of the SoTA kinematic planner, as it operates over a discrete version of the strain maps, resulting in various non-differentiable points along the commanded reference. Beyond resulting in higher shoulder accelerations, the SoTA planner also applied higher peak forces to the human arm in our experiments (Figure 3.5, bottom row), a result that can be attributable to ignoring human dynamics during planning.
The different modalities (i.e., an active and a passive subject) explored in this study can provide solutions for different stages of the therapy. After surgery or in early treatment after an injury, the patient is generally restricted in their mobility, and movements are carefully applied by the therapist who supports the arm and shoulder. We were able to control robot-led movements over long-horizon trajectories to increase the range of

motion of the shoulder while minimizing tendon strain in the rotator cuff to avoid re-injury (Figure 3.5). As treatment progresses and as the patient regains strength and mobility, patients become and are encouraged to be more active during the exercises, in which case navigation must be more responsive (Figs. 3.6-3.7-3.8).
In conclusion, our results demonstrate that embedding a high-fidelity representation of the human musculoskeletal system in the controller improves physical human-robot interaction in the context of automated yet responsive rehabilitation. In this study, we focused on rotator cuff strain monitoring and minimization, but other quantities could also be included (such as joint reaction forces or muscle activations themselves). Overall, the use of high-fidelity biomechanical models unlocks the design of robotic controllers that can monitor musculoskeletal quantities that even experienced therapists lack access to, potentially enhancing the utility of these systems in clinical rehabilitation. The next step is to evaluate the system on target users and patient groups.

3.6.4. Limitations and future work

We demonstrated and evaluated key aspects of BATON that provide confidence in its application for physiotherapy in the future. Nonetheless, we acknowledge that there are several limitations that require further investigation. First, we focused on the mobility of the glenohumeral joint alone, and planned human motions along two of its DoFs (PE and SE). In our setting, the motion along the third DoF (AR) of the human glenohumeral joint was not optimized, and the current value of AR determined the strain map on which the optimization took place. This simplification allowed us to validate the core concept, but spatial (3D) rehabilitation trajectories should be investigated for use in therapeutic applications. For example, the strain of the rotator cuff tendon depends on the relative position between the upper arm and the shoulder blade, and other movement DoFs, such as the scapula mobility or the human torso orientation, will also influence tissue loading in the shoulder.
Second, the human pose estimation that we employed did not include the complete human state and may rely on simplifying assumptions that are too restrictive for a wide range of physical therapy applications. While an accurate estimation of the subject's body motion was beyond the scope of this work, better methods to monitor human posture and movement could be integrated. Viable options could rely on motion capture techniques, which offer high precision but are difficult to deploy, while video camera systems or IMUs might be more practical but also less reliable. Exoskeletons could also provide better human kinematics with respect to the collaborative robotics arm employed in this study. In the future, we plan to explore how to extend BATON to also include additional DoFs, such as those related to the scapula, with an approach to measure the state of the scapula itself. However, apart from resulting in more complex computations, a challenging aspect of treating additional DoFs is delivering visual information that a patient/therapist can easily interpret.
A third limitation lies in the strain value delivered by the shoulder model we selected. While the model has been validated in previous studies for what concerns the accuracy of shoulder kinematics [102] and muscle activity [58, 59], a thorough validation of the estimated tendon behavior across the model's range of motion is necessary before we can apply robotic physical therapy for rehabilitation from rotator cuff tendon tears on a

patient population.
Lastly, we did not predict human behavior at planning time, opting instead for online adaptation to volitional actions and reflexive reactions. While BATON could be extended to include predictive models of human motion or neuromuscular dynamics(e.g., [155, 156]), they would provide only uncertain predictions that deviate from actual human behavior. The fundamental challenge of reliably predicting human decisions and behaviors remains. As such, we designed BATON to use actual behavior, necessary to ensure safety and effective performance, prioritizing real-time responsiveness to changes in human muscle activations estimated from measured movement and interaction forces. Future advances in biomechanical estimation could further enhance BATON's performance.

3.7. Conclusion

We presented BATON, a novel approach to physical human-robot interaction and its application for shoulder rotator cuff physiotherapy. Through biomechanics-aware trajectory optimization, BATON generates and executes rehabilitation trajectories applied to a human subject by a collaborative robot, leveraging a personalized musculoskeletal model. Mechanical tendon strain was monitored and minimized both during the early, passive stages of rehabilitation, and as muscle activation emerges when the subject responds, either actively or reactively, during therapy.
Unlike prior work in robotic-assisted physiotherapy, which does not directly monitor internal biomechanical metrics due to modeling and computational complexities, we have taken the first step to unlock the power of high-fidelity musculoskeletal models for real-time biomechanics-aware robotic controllers. Moreover, integral to BATON is an online alternative to sensor-based muscle activity estimations, such as surface EMG, which cannot measure underlying muscles like the rotator cuff.
While our primary focus is robotic rehabilitation, the ability of automated systems to estimate and respond to the changing internal states of the human body resonates far beyond physical therapy. BATON represents a foundational shift toward the next generation of human-centered assistive and augmentative robotics. We pass the "baton" to researchers and engineers in ergonomics, exoskeleton control, movement disorders, and beyond, where awareness of human biomechanics and subject/patient response is needed for robots to transform outcomes.

Appendix

Experimentally, commanding the Cartesian end-effector pose through equation (3.15) produces non-negligible tracking errors along the vertical axis if the human subject cannot support the weight of their own arm. Since full-arm support is typical of early rotator cuff rehabilitation [157], we implemented a gravity-compensation strategy modifying equation (3.15) to include the generalized torques $\mathbf{u}^*$ into ${}^{\text{base}}\bar{\mathbf{p}}_{\text{EE}}$. Let us define the end-effector Cartesian position as ${}^{\text{base}}\bar{\mathbf{p}}_{\text{EE}} = {}^{\text{base}}\bar{\mathbf{p}}_{\text{GH}} + {}^{\text{base}}\bar{\mathbf{R}}_{\text{EE}}\,{}^{\text{base}}\bar{\mathbf{p}}_{\text{GH}} = [p_x\ p_y\ p_z]^T$, similar to equation (3.15). From our biomechanics-aware trajectory optimization (Section 3.2.2), we obtain the human generalized torques needed to track such a

reference at every time instant. Thus, to deliver the required model-based torque along the *SE* DoF to the human subject, we added to the vertical EE reference a displacement δ_z:

$$\delta_z = \frac{u^{*,\mathrm{SE}}}{K_z L_{\mathrm{hum}} \sin(\hat{\mathrm{SE}})} \tag{3.17}$$

resulting in ${}^{\mathrm{base}}\tilde{\mathbf{p}}_{\mathrm{EE}} = [p_x,\ p_y,\ p_z + \delta_z]^T$

This approach is personalizable to each subject through their specific biomechanics and appears more robust to model inaccuracies than purely torque-based compensation methods (i.e. [158]). We show experimentally that this adjustment achieves the desired vertical EE setpoint, fully supporting the human arm without resorting to impractically high stiffness in the controller.

Figure 3.9 shows the Cartesian reference positions for the robot EE together with the actual ones across the duration of the experiment. In particular, the bottom row of the figure presents the tracking performance along the robot's vertical axis. It can be noted that the vertical reference $z_{ref} = p_z + \delta_z$, accounting for the optimized offset in equation (3.17), is not reached. However, this allows the Cartesian impedance controller to provide the correct support, such that the optimal end-effector height $z_{\mathrm{opt}} = p_z$ is tracked.

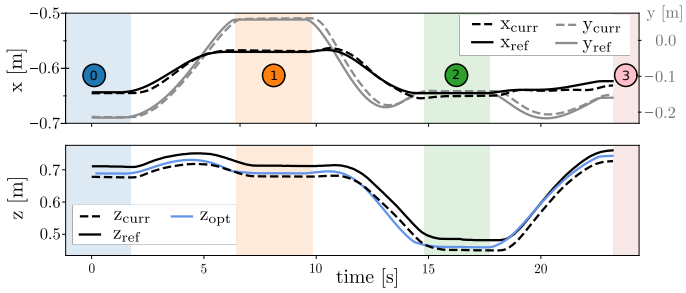

Figure 3.9: *Tracking performances of BATON's optimal trajectory during passive therapy, expressed into the Cartesian coordinates of the robotic end-effector* ${}^{base}\mathbf{p}_{EE}$. *The bottom graph highlights the effect of using the personalized musculoskeletal model to inform the gravity compensation of the human arm: the offset reference given to the controller (solid black) makes it so that the current position (dashed) better tracks the optimal one (solid blue).*

4

Real-time rotator-cuff tendon strain estimation

The previous chapter introduced BATON, a novel approach to robotic-assisted rehabilitation that tightly integrates a musculoskeletal model into robot control. By leveraging the RMR solver for online muscle activity estimation, BATON estimated target tendon strain online and generated safe, adaptive shoulder rehabilitation trajectories for our participant.

The present chapter elaborates further on our approach to estimating tendon strains by presenting in detail the concept of "active strain maps", which capture the effect of shoulder kinematics, external forces, and muscle activation on the resulting tendon strain. Such maps, which BATON used to guide rehabilitation in the previous chapter, can also serve as a visualization tool for human physiotherapists to help them understand how specific exercises affect their patients.

We demonstrated this by recording data from a participant performing common shoulder rehabilitation exercises against varying levels of robotic resistance. Inspecting how tendon strain evolved across these movements provides physiotherapists with a means to refine therapy and adapt exercises (whether delivered manually or robotically) ultimately enabling more personalized and effective care.

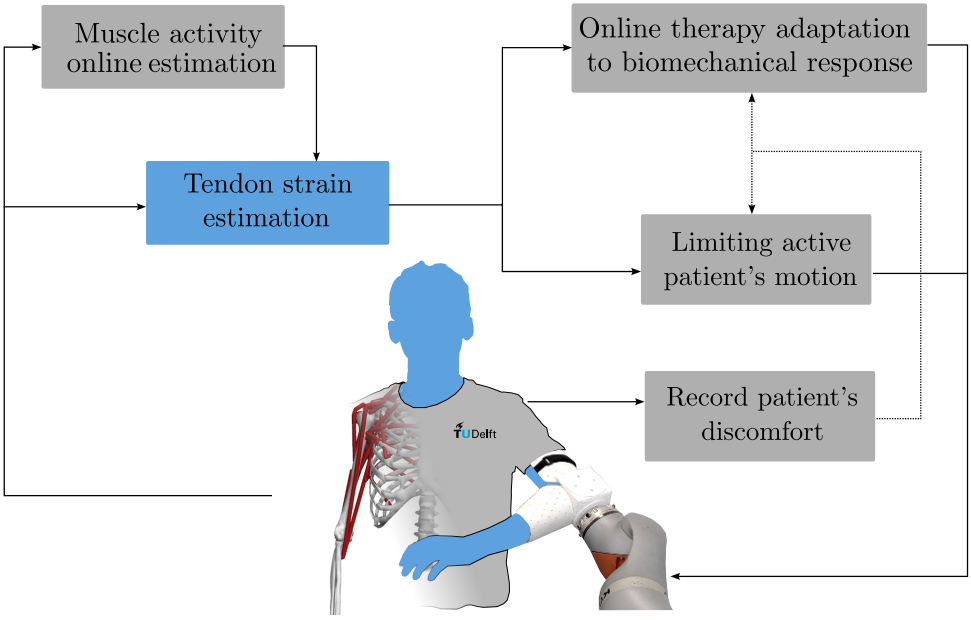

This chapter is based on I. Beck, **I. Belli**, L. Peternel, A. Seth, and J. M. Prendergast Real-Time Tendon Strain Estimation of Rotator-Cuff Muscles during Robotic-Assisted Rehabilitation IEEE-RAS 22nd International Conference on Humanoid Robots (Humanoids). IEEE, 2023. [147]. Finalist for the Best Paper Award.

4.1. Introduction

The shoulder complex is one of the most common sites for musculoskeletal disorders[159], with 70% of shoulder complaints attributed to disease or injury to the rotator cuff (RC) muscle-tendon units [160]. Tearing of the rotator cuff tendons results primarily in pain, loss of strength, and a restricted range of motion of the upper limb, which may affect the ability to perform work and/or activities of daily living [161]. The likelihood of RC tears also increases with age, with over 50% of people aged 60 or older suffering from one or more RC tears [162, 163]. In addition, repetitive overhead activities during work or competitive sports also increase the risk of RC tears [164–166].

Conventional treatment of RC tears requires intensive rehabilitation via physiotherapy. This rehabilitation may follow surgical repair in the case of severe tears, but is the standard of treatment for all RC injuries. Effective rehabilitation after surgery is necessary to prevent shoulder joint stiffness, regain range of motion and limit the risk of re-tearing [167]. Regardless of the treatment type, RC rehabilitation is time-intensive with significant improvement occurring after twelve weeks [168] and many patients are instructed to continue home exercises for as long as six months to a year [169, 170].

Due to the complexity of the shoulder joint and a lack of quantitative insight into the risks of re-tear, the exercises used in conventional RC rehabilitation are typically highly conservative and comprised of single-degree-of-freedom motions even when assisted by an experienced physiotherapist [123]. It is likely that this limited range of motion during rehabilitation may delay recovery, whereas moving through a larger range of motion could result in expedited recovery [171, 172]. Additionally, current rehabilitation techniques are demanding for physiotherapists who must physically manipulate many patients per day and who are limited to assisting one patient at a time [173]. Finally, treatment of the rotator cuff is often inadequate, leaving patients with persistent symptoms [174], inhibiting the patient's ability to work and perform daily activities in the long term.

Robotic systems can serve to address many of the limitations of conventional physiotherapy — providing a means for physiotherapists to treat more patients and to customize rehabilitation to a patient's specific needs. Several robots and exoskeletons exist for the rehabilitation of the upper limb [20, 175–177]. Robotic devices can provide high-intensity, repetitive exercises that target specific injuries/joints [178, 179]. Additionally, a robotic system can leverage its sensing capabilities to measure joint positions, velocities, accelerations, and torques [180] and provide force feedback. These metrics can be monitored objectively and reliably throughout a rehabilitation program to assess the patient's progress [173].

In the case of rehabilitation of RC tears, monitoring the subject's muscle tension is of interest to prevent (re-)tearing of the tendons. The nature of the RC as deeper (below superficial) muscles, combined with the soft-tissue structures of the glenohumeral (GH) joint, makes lab-based motion-capture and electromyography (EMG) measurements impractical for evaluating the function of RC muscles. Musculoskeletal modeling, in combination with robotic measurements, could provide a physiotherapist with quantitative information on the biomechanics of the shoulder complex and RC muscles. In our previous work, we developed a method to link musculoskeletal modeling to robotic rehabilitation by capturing the relationship between the tendon strains and shoulder

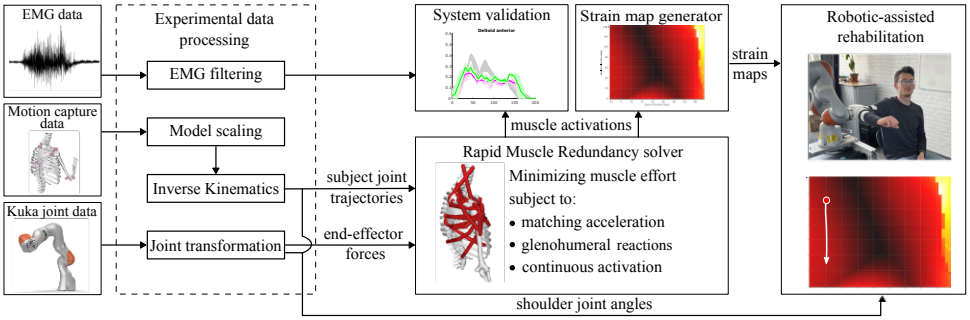

Figure 4.1: *Workflow for active robotic-assisted shoulder rehabilitation that provides quantitative insights into the risk of re-injury. The experimental data (left) is obtained via motion capture using optical markers or shoulder joint estimates during a human-robot interaction. The experimental data is processed to obtain the joint trajectories and external forces. The Rapid Muscle Redundancy solver (center) is used to estimate muscle activations from a musculoskeletal model for a given (measured) trajectory. The estimated muscle activations are validated against experimental EMG data and used to compute rotator cuff tendon strains. These strains are visualized in "strain maps" (bottom right) and are used to visualize the internal biomechanics of the shoulder during rehabilitation.*

state in the form of "strain maps" [115]. These strain maps are an intuitive representation of the RC tendon strains and can be used to minimize the risk of re-injury of the RC tendons while maximizing the range of motion. We recently implemented strain maps to enable a subject to perform robotic-assisted exercises while being protected from potentially dangerous poses [61] in real-time. A major limitation of these maps, however, is that they do not incorporate muscle activation into the muscle-tendon strain estimates. This restricts the use of these maps to passive motions (i.e., no external forces) in which the weight of the arm is supported entirely by the robot or a physical therapist, and thus limits the use of this tool to early-stage rehabilitation.

To address this limitation, we have developed a novel active strain maps approach that accounts for muscle contributions in various shoulder states (i.e., shoulder joint coordinates and external forces) as sensed by a collaborative robotic manipulator. We exploit our recently developed open-source Rapid Muscle Redundancy (RMR) solver [59] to estimate shoulder muscle activity. The proposed method then transforms these muscle activity estimates into our estimates of tendon strain for all of the RC tendons and creates active strain maps. Finally, we integrate these active strain maps with our robotic physical therapy (PTbot) system using the Kuka LBR iiwa 7 robotic arm. We demonstrate the use of these active maps with the collaborative robot during several conventional shoulder exercises.

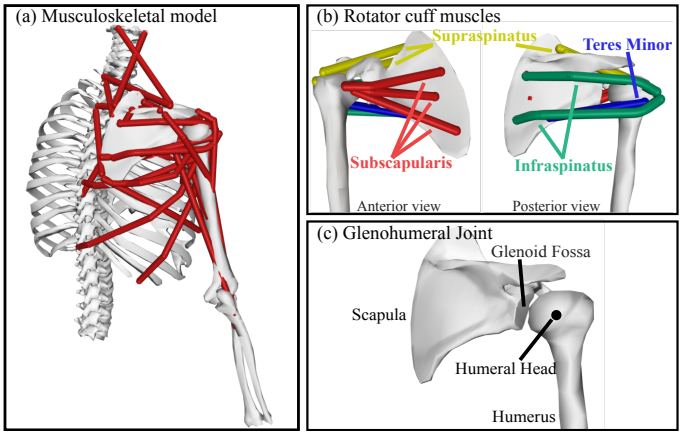

Figure 4.2: *An overview of the thoracoscapular model. (a) The musculoskeletal shoulder model, and in red, all the actuating muscles. (b) The rotator cuff complex, the four rotator cuff muscles are represented by eight muscle elements; Infraspinatus - green, supraspinatus - yellow, subscapularis - red, teres minor - blue. (c) The glenohumeral joint, detailing the relation between the humeral head and the glenoid fossa of the scapula.*

4.2. Methods

The methodology is divided into five subsections. Section 4.2.1 gives an overview of the musculoskeletal shoulder model that is leveraged to obtain information on the internal biomechanics of the shoulder during rehabilitation. Section 4.2.2 provides a summary of the muscle redundancy problem and how we solve for muscle activations, while considering active muscle dynamics and the GH joint stability. Section 4.2.3 shows the validation of the Python-RMR solver. Section 4.2.4 explains how our PTbot system is controlled to provide safe physiotherapy. Section 4.2.5 details the experimental design incorporating the active strain maps with the PTbot system. The proposed method is outlined as a block diagram in Figure 4.1.

4.2.1. Musculoskeletal Shoulder Model and Strain Maps

We leveraged the thoracoscapular musculoskeletal shoulder model [181] and the open-source simulation software OpenSim [182] to estimate rotator cuff muscle forces and tendon strains during shoulder exercises. This model, featuring $N = 7$ degrees of freedom (DoFs) and actuated by $N_m = 33$ muscle bundles, is designed to provide high fidelity and accurate representation of the shoulder complex by decoupling the movement of the humerus from the scapula. The complete shoulder kinematics are represented by 4 joints: the scapulothoracic [102], glenohumeral, acromioclavicular, and sternoclavicular joints. Of particular interest in our study are the 4 rotator cuff muscles: infraspinatus, subscapularis, supraspinatus, and teres minor, which span and stabilize the glenohumeral joint. The glenohumeral joint is represented as a three

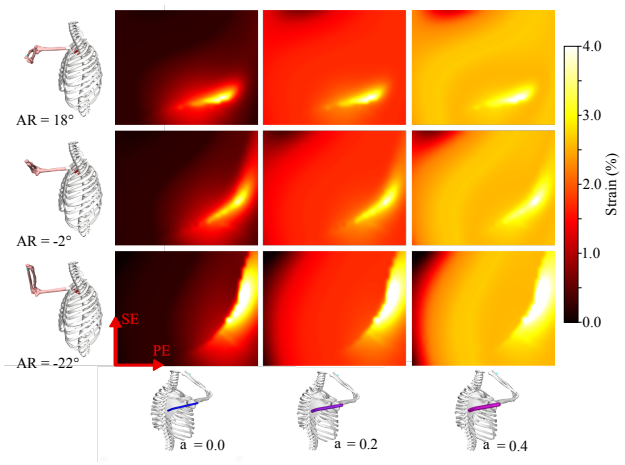

Figure 4.3: *Interpreting tendon strain from the visualization of strain maps. This example considers the infraspinatus inferior; it includes nine strain maps for a combination of three different axial rotations and three muscle activation levels. The individual maps include the plane of elevation and shoulder elevation on the x- and y-axis, respectively. The columns represent the strain maps for increasing muscle activation α of the infraspinatus inferior (α = 0.0, α = 0.2, and α = 0.4). The rows show strain maps for selected axial rotation angles (AR = -22°, AR = -2°, and AR = 18°).*

DoFs gimbal joint (Figure 4.2), and the rotator cuffs primarily affect the direction of the glenohumeral joint reaction force. The strain in the rotator cuff tendons is a measure of (re-)tearing risk [61, 115]. Muscle strains are dependent on the shoulder state and muscle activations, which must be accurately and reliably estimated.
Computing the RC tendon strains from the musculoskeletal model directly is computationally expensive and only provides individual strain values in numerical form, which would be difficult to interpret on its own during rehabilitation. Instead, we chose to enumerate RC muscle strains into a simple yet comprehensive and intuitive view of the possible strain landscape we call “strain map” [61, 115] that enables the physiotherapist to see where they are going in terms of increasing or decreasing strain. Strain maps provide the user with a high-resolution visualization of the tendon strain in the joint space. Different combinations of strains can be utilized depending on the patient’s injury and progress during therapy. For example, the strain maps of all rotator cuff muscle elements can be combined into a single strain map by taking the maximum strain for each pose. A strain map may also be visualized for a single muscle (element) to ensure that strains are safe during rehabilitation of a specific injury with a known muscle-tendon location (i.e., a surgical repair of the supraspinatus).
Since the rotator cuff muscles only span the GH joint, we consider the 3 shoulder DoFs as the state $\boldsymbol{\theta}_S$ of our model, consisting of the following Euler angles; plane of elevation (PE), shoulder elevation (SE), and axial rotation (AR):

$$\boldsymbol{\theta}_S = [\ \mathrm{PE},\ \mathrm{SE},\ \mathrm{AR}\ ]\ . \tag{4.1}$$

The range of shoulder DoFs was constrained within the physiological limits of $-20° \leq PE \leq 160°$, $0° \leq SE \leq 144°$, and $-90° \leq AR \leq 100°$. Rotator cuff tendon strain estimations were pre-computed from the musculoskeletal model for each combination of $\boldsymbol{\theta}_{\mathrm{S}}$ at 4° intervals and muscle activation in 0.005 activation level intervals.
To visualize the four-dimensional space for a physiotherapist, the strain maps were divided into two-dimensional layers/maps with the plane of elevation and shoulder elevation on the x- and y-axis, respectively. The axial rotation and muscle activation are fixed for every two-dimensional map. For real-time visualization purposes, the 2D map is updated if axial rotation of muscle activation is changed, as shown in Figure 4.3.

4.2.2. Muscle Activation Estimation

A fundamental estimate used by our strain calculation is the muscle activation level required during shoulder movements. Due to the redundancy in muscle elements compared to the DoFs of the model, there are infinite solutions for combinations of muscle activations for any given movement/loading condition. To determine a unique solution, we turn to optimization and employ the rapid muscle redundancy (RMR) solver [59]. RMR solves for individual muscle activations that generate observed estimates of joint accelerations such that the sum of muscle effort is minimized. The model computes muscle forces that account for musculotendon passive forces, muscle activation dynamics and glenohumeral joint stability to estimate physiologically realistic muscle activations in real-time.
The solver's inputs are biomechanical joint trajectories and external forces/torques acting on the subject (e.g. at the interface between robot and subject). Muscle activations are estimated by solving a nonlinear optimization problem at each time instant t_k of the human motion, minimizing the biological muscle effort expressed as a weighted sum of squared muscle activations $\mathbf{a}_k \in \mathbb{R}^{N_m}$:

$$\min_{\mathbf{a}(\cdot),\mathbf{c}(\cdot)} \quad \sum_{i=1}^{N_{\mathrm{m}}} w_i a_{i,k}^2 + \sum_{j=1}^{N} \mu_j c_{j,k}^2 \tag{4.2a}$$

$$\text{subject to:} \quad \mathbf{A}_{\mathrm{acc},k} \begin{bmatrix} \mathbf{a}_k \\ \mathbf{c}_k \end{bmatrix} = \ddot{\boldsymbol{\theta}}_{\mathrm{exp},k} - \ddot{\boldsymbol{\theta}}_{\mathrm{pass},k} \qquad \text{acc. matching} \tag{4.2b}$$

$$\mathbf{F}_{\mathrm{GH}} = \mathbf{A}_{\mathrm{F},k} \begin{bmatrix} \mathbf{a}_k \\ \mathbf{c}_k \end{bmatrix} + \mathbf{F}_{0,k} \in \mathscr{C} \qquad \text{GH stability} \tag{4.2c}$$

$$\forall\, i \in N_{\mathrm{m}} : \quad l_{i,k} \leq a_{i,k} \leq u_{i,k} \qquad \text{activation dyn.} \tag{4.2d}$$

Because the model is a simplified version of reality, reserve actuators were added to the model to account for model discrepancies like muscles being slack or too stretched (as recommended in [100]) and their control levels $\mathbf{c}_k \in \mathbb{R}^N$ were also included in the objective function equation (4.2a). The relative weights (penalties) w_i and μ_j were set to 1 and 10, respectively, to promote the engagement of the muscles over the reserve actuators, which capture modeling errors.

Constraints are used to guide the solver towards physiologically realistic solutions. First, we ensure that the simulated joint accelerations match the experimental data ($\ddot{\boldsymbol{\theta}}_{\text{exp},k} \in \mathbb{R}^N$) for all the DoFs, by means of a linear equality constraint in equation (4.2b) where we account purely for the active muscle contribution, subtracting from $\ddot{\boldsymbol{\theta}}_{\text{exp},k}$ the influences of gravity and passive muscle forces at the given shoulder pose, collectively lumped in $\ddot{\boldsymbol{\theta}}_{\text{pass},k} \in \mathbb{R}^N$. Element $A_{\text{acc},k}(j,i)$ of $\mathbf{A}_{\text{acc},k} \in \mathbb{R}^{N\times(N_m+N)}$ represents the influence of a single activation of actuator i on the acceleration of coordinate j. Secondly, a stability constraint in equation (4.2c) is enforced on the joint reaction force $\mathbf{F}_{\text{GH}}$ at the glenohumeral joint. The direction of the $\mathbf{F}_{\text{GH}}$ was constrained to be within a circular approximation of the glenoid fossa $\mathscr{C}$, to ensure the stability of the GH joint. The last set of constraints in equation (4.2d) ensures physiological activation and deactivation rates, as muscle's force generation and relaxation are dependent on the calcium ion concentrations. The dimensionless muscle activations can range between 0 and 1, where a value of 0 indicates no contraction of the fiber (no activation), and a value of 1 represents maximum muscle contraction and thus force generation (full activation).

4.2.3. Implementation of a streaming Python-RMR solver

To facilitate the integration of the RMR solver with the proposed active strain map method and a robotic system, the original implementation of the solver in MATLAB [59] was re-implemented in Python with a few important distinctions. The solver makes use of the Sequential Least Squares Programming (SLSQP) method from the Scipy.minimize library [183]. In addition, the Python implementation accommodates real-time applications by allowing for streaming pose and external forces data, which enabled us to take measurements from the robotic manipulator and compute muscle activity at approximately 5Hz such that the strain maps can be updated during an exercise based on estimated muscle activations. We tested the Python-RMR implementation against the original to verify our muscle activation estimates.

Muscle activations estimated with the RMR solver in MATLAB and Python were compared to the experimental EMG data. The activation was averaged over the three repetitions for each task. The muscle activation estimations were compared during three motions: abduction, flexion, and shrugging for both an unloaded and 2 kg load condition. To capture differences across the dataset, mean absolute errors (MAE) were computed. The MAE values for MATLAB and Python against EMG-based activations were found to be ≤ 0.17, with most values ≤ 0.10. For the majority of the activation estimations, the errors against EMG data were the same in the MATLAB and Python implementations. The differences in MAE values between MATLAB and Python were ≤ 0.02. Figure 4.4 included the mean and standard deviation of muscle activations for selected muscles during a loaded flexion task. The results from the experimental EMG data and the MATLAB and Python RMR implementation demonstrate that the Python-RMR solver produces results nearly identical to the MATLAB implementation, with only minor differences. Finally, we computed the average frame processing frequency of the Python-RMR solver to be approximately 5 Hz.

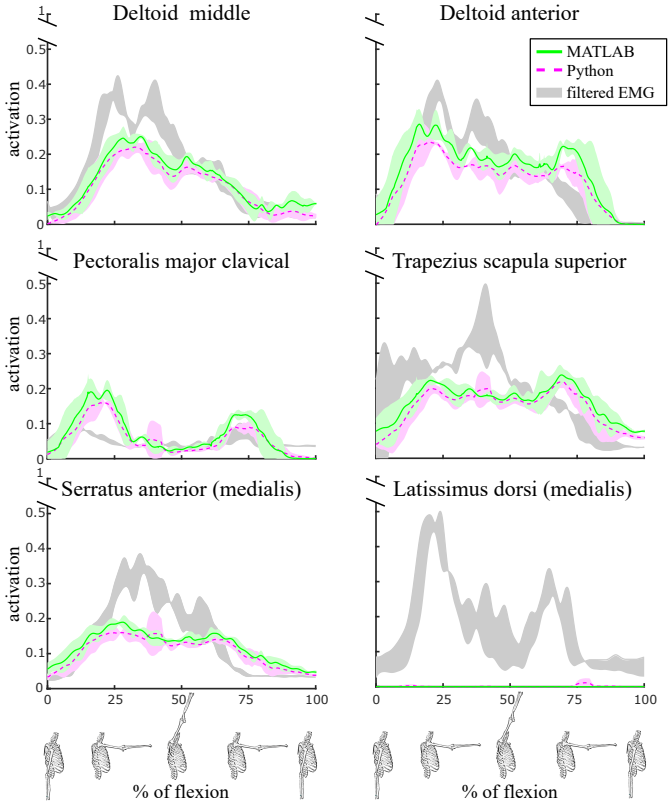

Figure 4.4: *Muscle activation of selected muscles during a loaded flexion motion. The estimations are shown for MATLAB (green) and Python (pink) and are displayed as the mean over three repetitions of the motion (bold lines), with a shaded* ±1 *SD. Additionally, the muscle activations from filtered experimental EMG data are displayed in gray (*±1 *SD).*

4.2.4. Robotic Control and Interaction

The strain maps need input about the shoulder state vector $\boldsymbol{\theta}_S$ and any external forces/torques acting on the subject during the therapy to retrieve the internal biomechanics of the shoulder joint. We leveraged the KUKA LBR iiwa, a 7-DoFs industrial collaborative robot certified for safe physical human-robot interaction. A 3D-printed plastic brace is used as a rigid interface between the robot and the human arm, and the subject was instructed to minimize the displacement of their shoulder during the interaction with the robot (Figure 4.5). In such a way, the shoulder state $\boldsymbol{\theta}_S$ and its time derivative $\dot{\boldsymbol{\theta}}_S$ can be uniquely obtained from the current pose of the robot's end-effector, while the other coordinates are kept fixed (see Section 3.3.2). The potential discomfort to the subject is reduced by leveraging gravity compensation for both the robot and the arm brace, and interaction forces at the human-robot interface

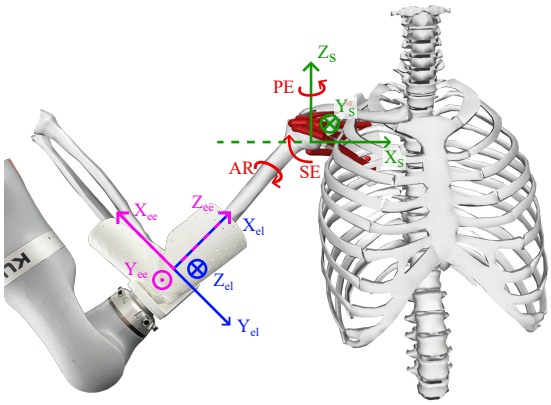

Figure 4.5: *An overview of the coordinate systems of the setup: The shoulder frame has its origin in the center of the glenohumeral joint (green). The glenohumeral joint DoFs (PE, SE, and AR) are shown in red. The elbow frame (blue) has an origin in the center of the elbow. The elbow frame origin coincides with that of the robot's end-effector frame (pink).*

are estimated from the joint torques [184]. Data from the torque sensors and the joint angles were collected at 200Hz, and subsequently used as input for the RMR solver for the muscle activation and tendon strain estimations. During our experiments, the robot was controlled with a Cartesian impedance controller:

$$\mathbf{F}_{\text{imp}} = \mathbf{K}(\bar{\mathbf{x}}_{\text{ee}} - \mathbf{x}_{\text{ee}}) + \mathbf{D}(\dot{\bar{\mathbf{x}}}_{\text{ee}} - \dot{\mathbf{x}}_{\text{ee}}) \,. \tag{4.3}$$

where $\mathbf{F}_{\text{imp}} \in \mathbb{R}^6$ is the vector of interaction forces and torques at the end-effector. The robot's end-effector reference and actual pose are $\bar{\mathbf{x}}_{\text{ee}} \in \mathbb{R}^6$ and $\mathbf{x}_{\text{ee}} \in \mathbb{R}^6$, respectively. $\mathbf{K}, \mathbf{D} \in \mathbb{R}^{6\times6}$ are the desired stiffness and damping matrices in Cartesian space. The positional and rotational stiffness of $\mathbf{K}$ are prescribed, and $\mathbf{D}$ is obtained using the double diagonalization design method [185]. By setting different positions of $\bar{\mathbf{x}}_{\text{ee}}$ at the beginning of the experiment, the robot mimics the effect of elastic bands during rehabilitation exercises (Figure 4.6).

4.2.5. Experimental setup

To test the integration of the muscle redundancy solver with the PTbot system, we conducted an experiment to determine RC tendon strain during four different exercises. One healthy individual was considered as a subject for the following experiments. They were seated in a normal chair and, after the robot was moved to its initial position, wore the custom arm brace. The subject performed four different movements based on common rehabilitation exercises; forward flexion, extension, abduction, and adduction (see Figure 4.6). They were free to move in space but instructed to make 1-DoF movements. For forward flexion and extension, axial rotation and plane of elevation

were kept at around −90° and 90°, respectively. In the case of abduction and adduction, axial rotation and plane of elevation were kept at around 0°. To ensure the subject's safety, their movement was restricted within a limited ROM. For all cases, shoulder elevation was instructed to be limited between 30° and 80°.
The Cartesian impedance controller of the robot was leveraged to mimic elastic bands commonly used in rehabilitation exercises. By increasing the positional stiffness, a larger external force acted on the elbow; thus, more effort was required from the subject to move along the trajectory. The four movements were all executed for three different positional stiffness values: K_{pos} = 10, 30, and 50 N/m. The rotational stiffness was kept constant throughout the experiments. The stiffness value K_{rot} was set to 50 Nm/rad.
To allow for the estimation of shoulder joint angles based on the robot joint angles, the subject was instructed not to move their torso. During the experiments with the robot, we only obtained information on the GH joint coordinates and thus assumed that the scapula was not moving during the muscle activation estimation. The shoulder state and external force data were post-processed to estimate the muscle activations and the tendon strains and generate the strain maps.
All the computations were run on an HP ZBook Studio G3 with an Intel i7-6700HQ processor and 8GB RAM.

4.3. Results

In this section, we show the resulting muscle activations and tendon strains from the conducted experiments with a robotic device for a set of motions and show the effect of external forces on the muscle activations and, subsequently, the tendon strains (4.3.1). Finally, we show the resulting strain maps, which now represent active model-based tendon strain estimations during human-robot interaction in 4.3.2.

4.3.1. Muscle activation and tendon strain estimation

The resulting muscle activation and tendon strains for the experiments performed with the Kuka LBR iiwa robotic arm are presented below. Two of the motions are shown, each with a different rotator cuff muscle highlighted. Figure 4.7 shows the resulting muscle activation and tendon strain of the infraspinatus inferior and infraspinatus superior during extension. A single extension was performed for three different robot stiffness values (K = 10 N/m, K = 30 N/m, and K = 50 N/m). By increasing the stiffness, the magnitude of the external force also increases. During extension, the infraspinatus superior is not activated for any of the modalities. The infraspinatus inferior is recruited less when the stiffness is increased, and the infraspinatus inferior tendon strain also reduces with higher stiffness. A peak at the end of the motion with medium stiffness can be seen in the infraspinatus inferior muscle activation and tendon strain. The supraspinatus muscle bundles are highlighted during a single abduction motion, again for all three stiffness values (Figure 4.8). Both the supraspinatus anterior and posterior bundles were activated during the abduction task. The estimated muscle activation increased with a higher positional stiffness of the robot, which also resulted in a higher tendon strain.

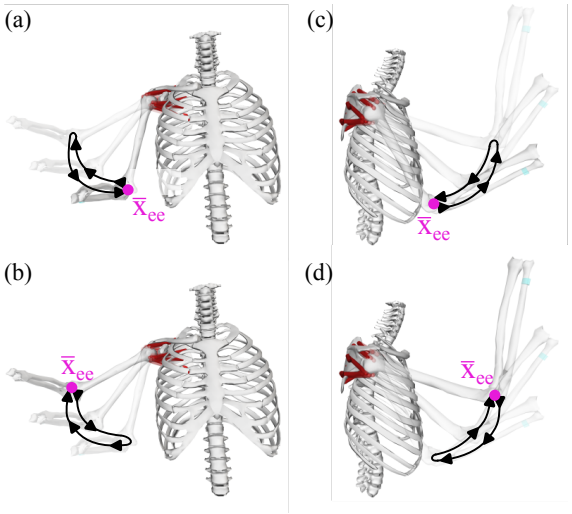

Figure 4.6: *The four motions as performed during the experiment with (a) abduction, (b) adduction. (c) forward flexion, (d) extension. The start position of each motion was set as the reference pose for the robot $\bar{x}_{ee}$ (pink), and the subject made one continuous motion as indicated in black. The arrows indicate the direction of the motion.*

4.3.2. Strain maps during active rotator cuff rehabilitation

Using strain maps, we visualized the tendon strains during a single abduction motion with medium positional robot stiffness. A subset of the resulting strain maps is shown in Figure 4.9. They represent the strain distribution for five instances in time throughout the motion. An example of strain maps is shown for subscapularis inferior (SSCI), infraspinatus superior (ISS), and the maximum strain of all rotator cuff muscles. These maps exhibit varying shapes as the shoulder pose and muscle activation levels change throughout the motion, resulting in a shift of low and high-strain zones. Notably, the two individual muscle elements experience high strains at different shoulder poses.

4.4. Discussion

In this study, we developed a method for the real-time consideration of external forces during robotic shoulder physical therapy. This is a significant step beyond our previous contributions, which could not account for the RC tendon strains induced due to muscle activity and instead assumed the patient was always passive during the exercises.

The inclusion of muscle activation in the tendon strain estimator provides more physiologically realistic estimates, allowing for a larger range of rehabilitation applications. This is particularly important in later stages of rehabilitation where a patient may begin to apply greater forces (resistance training), but may still be important in earlier stages if for example the patient is simply supporting the weight of their own arm. In general, the stabilizing role of the RC muscles makes their activation (and thus

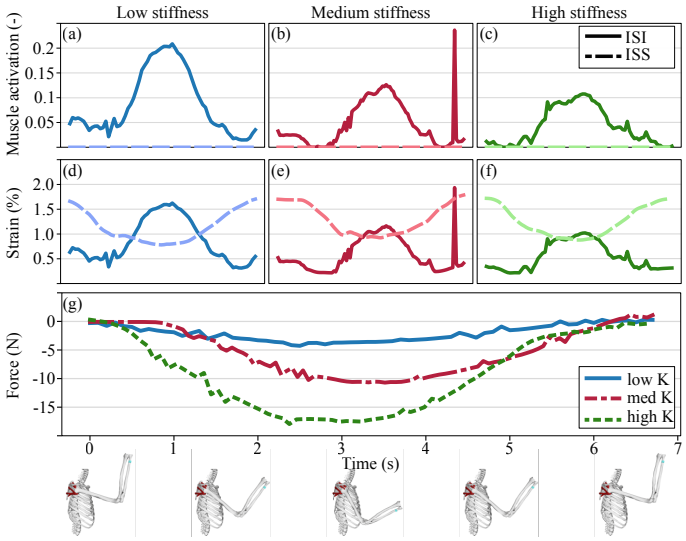

Figure 4.7: *An example of the shoulder extension movement and the estimated muscle strain and tendon activation of the infraspinatus inferior (ISI) and infraspinatus superior (ISS). The muscle activations (a-c) and tendon strains (d-f) are shown for all three robot stiffness values. In (g), the magnitude of the external force at the elbow is shown for the three different conditions.*

the tendon strains) quite difficult to predict. Providing accurate estimates of how the activations and strains will be induced, even during fully assisted motions, is a critical step for safe robotic assistance during rehabilitation.

To prevent injury and improve rehabilitation outcomes, it is critical for a robotic PT to understand how a patient should or should not move during rehabilitation. In the case of RC injuries, this requires insight into the RC muscle recruitment patterns and the impact these have on RC strains. The active strain maps presented here are unique in their ability to make these estimates in real-time and offer opportunities to develop new novel rehabilitation exercises tailored to the patient. In the simplest case of early-stage rehabilitation (where avoiding re-tear following surgery is crucial), a physiotherapist can choose the robot's stiffness and trajectory such that the strains in the damaged muscles are below a certain threshold. For example, in the case of an injured infraspinatus, a higher stiffness during extension was observed to reduce the strains. Conversely, during the later stages of rehabilitation, a physiotherapist may desire to prioritize strengthening a specific muscle or recruit compensatory muscles to compensate for an injured one. The combination of a customizable robotic device and the solver presents an intriguing opportunity to design and evaluate novel rehabilitation exercises for various rotator cuff disorders. This approach allows for exploring and establishing tailored rehabilitation protocols, providing valuable insights for improving patient recovery and outcomes.

The addition of muscle activation also adds substantial complexity to the strain mapping problem. While passive strain maps can be fully known beforehand, active maps rely on

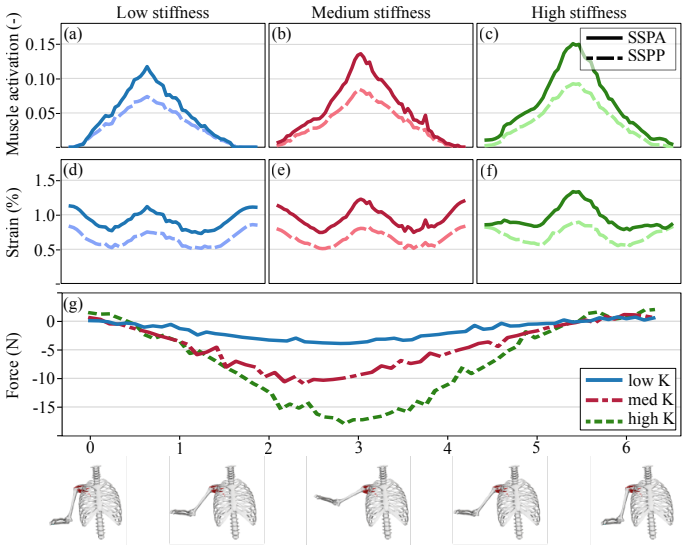

Figure 4.8: *An example of the shoulder abduction movement and the estimated muscle strain and tendon activation of the supraspinatus anterior (SSPA) and supraspinatus posterior (SSPP). The muscle activations (a-c) and tendon strains (d-f) are shown for all three robot stiffness values. In (g), the magnitude of the external force at the elbow is shown for the three different conditions.*

external forces applied by the patient, which are largely unpredictable. Real-time map updates based on these external forces are thus required to correctly reflect the tendon strains during active exercises. It is thus uncertain how the strain space will evolve in each subsequent time step. For clinical implementation, further investigation is required to enable real-time updates of active strain maps while incorporating the previously identified safe and unsafe strain zones [61]. We expect to leverage tools such as model predictive control in future studies as we incorporate this complexity into our PTbot system. Additionally, the 5 Hz update speed of the Python-RMR solver is fast enough for the speed of most rehabilitation exercises, since human movement frequencies are typically within the range of 0 to 10 Hz [186]. We will improve this in the future via dedicated hardware and software.

Finally, although the object of this work is focused on robotic rehabilitation, the results also demonstrate the potential for this system to provide important insights to a human PT during physical therapy exercises. In the shoulder extension experiment for example, we observed a decreased engagement of the infraspinatus inferior as a response to the larger external forces applied to the subject (Figure 4.7). Notably, the tendon strains are highly dependent on changes in shoulder pose even under low muscle activity, as observed in the infraspinatus superior.

An opposing muscle response to external forces was observed in Figure 4.8, where the supraspinatus superior and supraspinatus inferior were activated more for higher forces. While the rotator cuff muscles are the primary stabilizers of the glenohumeral

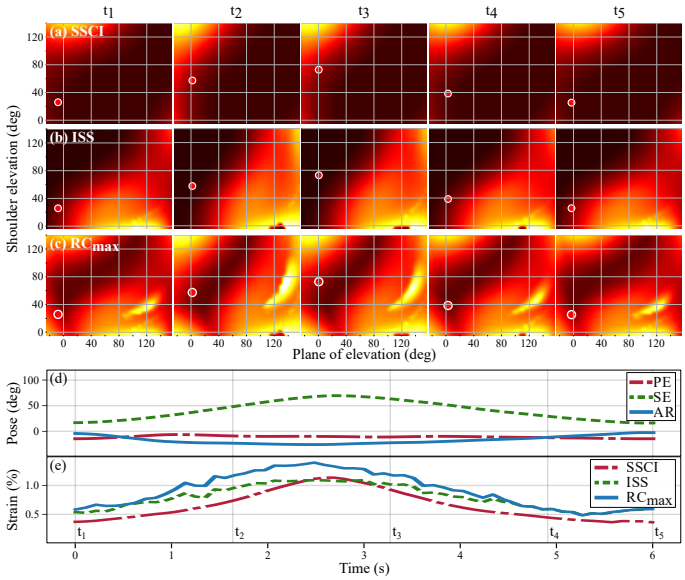

Figure 4.9: *An example of strain maps visualizing tendon strain at five instances ($t_1 - t_5$) during a single abduction motion with a medium stiffness (K=30 N/m). The strain maps are shown for the a) subscapularis inferior (SSCI), b) infraspinatus superior (ISS), and c) the maximum strain of all rotator cuff muscles. For the complete movement, d) the shoulder pose and e) selected rotator cuff tendon strains are shown.*

joint, they also assist in the motions of the shoulder joint. The recruitment shown in Figure 4.8 of the supraspinatus during abduction aligns with our intuition around shoulder biomechanics for the changing external forces. Because the supraspinatus assists the deltoids in abduction, increased recruitment of both muscles was expected for increased antagonistic external forces. This recruitment can be seen clearly in Figure 4.8 in response to the increased stiffness of the robot.

The infraspinatus is known as the primary external rotator of the humeral head, so engagement during the extension motion, with a 90° external rotation of the humerus, is in line with these expectations. In contrast, we expected larger external forces to destabilize the glenohumeral joint and, thus, a higher engagement of the rotator cuff muscles, which was not the case for the infraspinatus inferior during extension. This may be explained by an observed increase in latissimus dorsi muscle activation, which helps depress the humerus inside the glenohumeral joint, and requires less stabilizing efforts from the infraspinatus.

These observations demonstrate the ability of this system to provide improved understanding not merely for our PTbot but also clear and useful insights for physical therapists as we begin to link musculoskeletal modeling with real-time sensing and control.

4.5. Conclusion

In this work, we developed active strain maps, a novel functional representation of rotator cuff tendon strains in response to pose, movements and external forces on the arm. These active maps are a significant step forward from our earlier passive strain maps, as they offer the ability to improve the accuracy of estimated RC tendon strains by considering the effect of muscle activation during motions and under load. We demonstrate the use of these active strain maps via our robotic physical therapy system PTbot – conducting four conventional RC therapy exercises under varying loading conditions. By connecting this active strain map framework with our PTbot system, we can provide physiotherapists with safe and unsafe ROM during rehabilitation, and the ability to visualize quantitative information on tendon strain during assisted or unassisted and even resisted exercises. These active maps can be generated for every muscle element separately or can be used to represent the maximum strain experienced by all the rotator cuff muscles. Using this tool, we expect to enable improved robotic perception of the RC tendon strains during robotic rehabilitation for use in the planning of custom exercises and for protecting patients during motions. Finally, strain maps paired with our PTbot system provide a novel and useful feedback tool for assisting human physiotherapists, who can receive real-time visual feedback of RC tendon strain during rehabilitation exercises and can thus leverage the different active strain maps and robot control settings to deliver therapy for specific injuries over a large ROM while reducing the risk of re-injury.

5

Safely limiting patient motion through shared control

The previous two chapters demonstrated how biomechanics-aware robots can optimize shoulder rehabilitation at opposite ends of the therapeutic spectrum: Chapter 3 introduced BATON for adaptive guidance in early passive therapy, while Chapter 4 applied the same principles to monitor fully autonomous patient movement through tendon strain maps.

Between these two extrema lies a critical stage of rehabilitation: assisting patients who are encouraged to perform active movements, while being still at risk of re-injury. Many rehabilitation protocols call for early mobilization, which can be harmful if performed prematurely.

To address this gap, this Chapter proposes an intermediate approach where control of the movement is shared between human and robot. Patient movement proceeds freely as long as it remains within safe, low-strain regions of the shoulder's range of motion. When unsafe movement is predicted, the robot intervenes via variable impedance control, resisting or redirecting motion to prevent excessive tendon strain.

We evaluated this strategy in experiments with a healthy participant simulating patient behavior, testing alternative modules for limiting unsafe motions. These results lay the groundwork for broader user studies to identify the most effective human–robot interaction strategy for this intermediate rehabilitation stage.

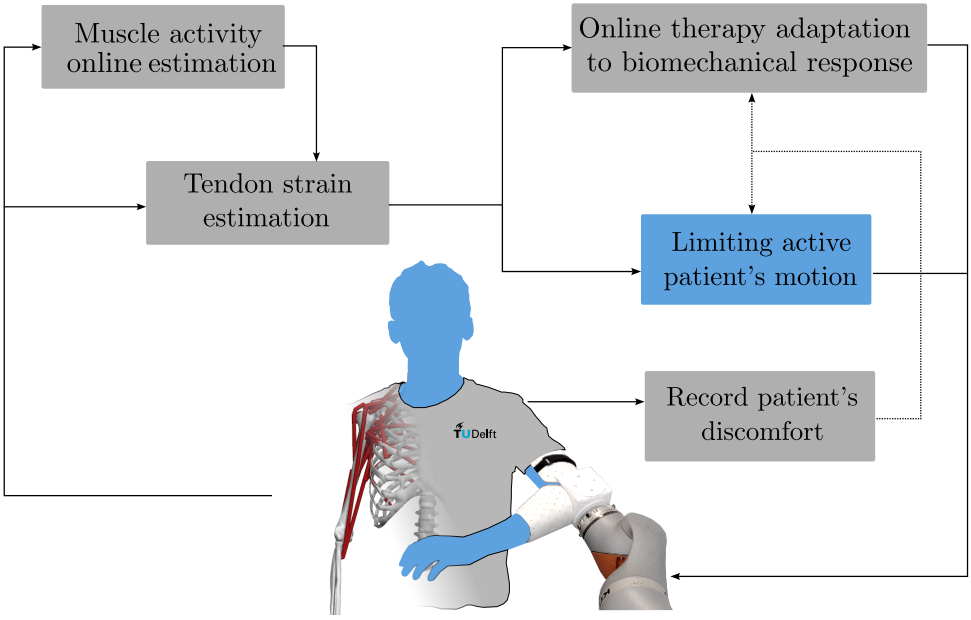

This chapter is based on **I. Belli**, Y. Hu, J.M. Prendergast, D. Abbink, A. Seth, and L. Peternel. A Shared Control Approach to Safely Limiting Patient Motion Based on Tendon Strain During Robotic-Assisted Shoulder Rehabilitation. 2025 International Conference On Rehabilitation Robotics (ICORR), 1071-1077. [187].

5.1. 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 the 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 [48].

Traditional therapy in treating injuries to the shoulder involves sessions with human physiotherapists. Nevertheless, re-injuries are common [55], 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 [128], as well as learning these movements from expert human demonstrations [129]. These strategies are excellent in automating the movements of human physiotherapists. 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 [130]. Most of the existing studies employ models offline, for example, to quantify assistance needed by a human operator [40, 131, 132] or reduce human metabolic cost in walking-assistive devices [41, 133, 134]. 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 [61, 140, 143, 144].

Our previous work [144] 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 [60], which enabled an online 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 [188], 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 [189, 190], 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 [61], 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

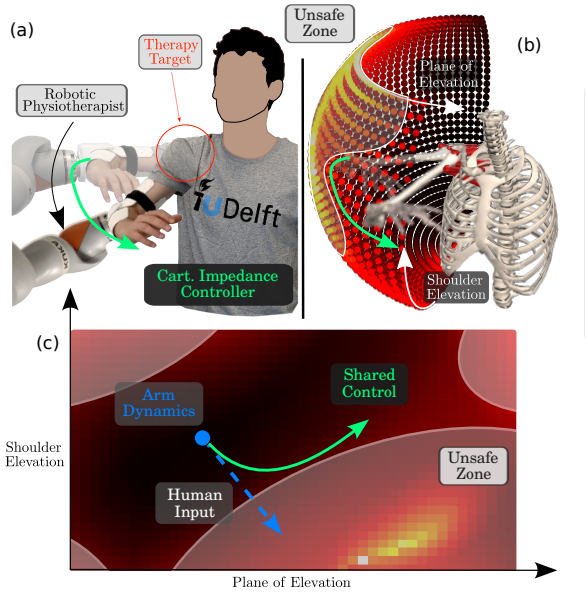

Figure 5.1: *An overview of the proposed robot-assisted shoulder physiotherapy system. (a) The robot end effector is connected to the human elbow with a customized brace, allowing the robot to limit human motion when necessary. (b) Shoulder strain map is pre-calculated from a high-fidelity musculoskeletal 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.*

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

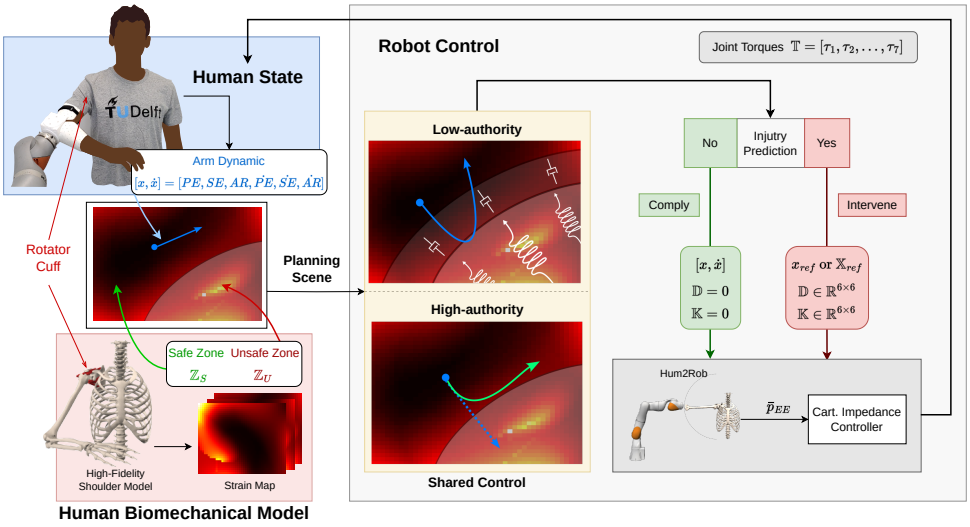

Figure 5.2: *Diagram of the proposed real-time system for safely limiting human motion during rehabilitation, on the basis of a high-fidelity OpenSim shoulder model. The two shared control modules receive the current state of the human shoulder and their corresponding strain maps to identify dangerous 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 complementary and can be selected based on the therapy requirements and user preferences.*

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.

5.2. Methods

Figure 5.2 illustrates the proposed system with key elements. The human modeling and strain map computation elements 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.

5.2.1. 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 [58] developed in OpenSim [31, 99]. 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:

$$\boldsymbol{x} = [\boldsymbol{\theta}_{\mathrm{S}},\ \dot{\boldsymbol{\theta}}_{\mathrm{S}}], \quad \text{with} \quad \boldsymbol{\theta}_{\mathrm{S}} = [PE, SE, AR] \tag{5.1}$$

where PE, SE, and AR are defined as the Y-X'-Y" sequence of intrinsic Euler angles in the fixed shoulder reference frame [58], and $\dot{PE}, \dot{SE}$ and $\dot{AR}$ are their corresponding derivatives with respect to time. Similar to our previous work [60, 144], 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 [148] 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:

$$\dot{\boldsymbol{x}} = f_{\mathrm{FD}}(\boldsymbol{x}, \boldsymbol{u}) \tag{5.2}$$

$$\boldsymbol{u} = f_{\mathrm{ID}}(\boldsymbol{x}, \dot{\boldsymbol{x}}) \tag{5.3}$$

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 [60] and brings the benefit of accounting for subject-specific parameters both regarding tendon behavior and human inertial properties.

5.2.2. 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 (Figure 5.2). As such, the estimation of the current human state $\hat{\boldsymbol{x}}_t = [\hat{\boldsymbol{\theta}}_{S,t}, \hat{\dot{\boldsymbol{\theta}}}_{S,t}]$ can be achieved based on the Cartesian pose $\boldsymbol{x}_{EE} \in \mathbb{R}^6$ of the end-effector (EE), assuming a fixed orientation of the human torso and negligible movement of the scapula. More details on this derivation can be found in Section 3.3.2 of this thesis.

5.2.3. Robot Control

The interaction between our rehabilitation robot and the human subject is shaped with an impedance controller [191] 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 $\boldsymbol{F}_{EE} \in \mathbb{R}^6$ can be defined as:

$$\mathbf{F}_{EE} = \mathbf{K}_t(\overline{\mathbf{x}}_{EE} - \mathbf{x}_{EE}) - \mathbf{D}_t\dot{\mathbf{x}}_{EE} \tag{5.4}$$

where the end-effector reference and actual pose are $\overline{\mathbf{x}}_{EE}$ and $\mathbf{x}_{EE}$, respectively, and $\mathbf{K}_t, \mathbf{D}_t \in \mathbb{R}^{6\times6}$ are the desired, time-dependent stiffness and damping matrices in Cartesian space. The selection of $\overline{\mathbf{x}}_{EE}$, $\mathbf{K}_t$, 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.

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 the estimated human pose $\hat{\boldsymbol{\theta}}_{S,t}$, 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 $\hat{\boldsymbol{\theta}}_{S,t}$ 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 $\hat{\dot{\boldsymbol{\theta}}}_{S,t}$. If $\hat{\dot{\boldsymbol{\theta}}}_{S,t}$ 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 $\mathbf{K}_{low}$ associated with damping exceeding the critical one $\mathbf{D}_t = 2r\sqrt{\mathbf{K}_{low}}$ (where we employ the ratio $r > 1$). In this way, the subject receives haptic information for safer direction of movement;

- if $\hat{\boldsymbol{\theta}}_{\mathrm{S},t}$ 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 [61]).

Algorithm 1 Low-Authority Shared Control

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

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 horizon T, to find the estimated future subject's trajectory $\hat{\boldsymbol{X}} = [\hat{\boldsymbol{x}}_{t+1}, \hat{\boldsymbol{x}}_{t+2}, ..., \hat{\boldsymbol{x}}_{t+T}]$. 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 5.2.3. If, instead, the current movement leads the subject to any unsafe zones, the algorithm determines an alternative trajectory $\boldsymbol{X}$ 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 $\mathbf{K}_t = \mathbf{K}_{\mathrm{high}}$, $\mathbf{D}_t = 2\sqrt{\mathbf{K}_{\mathrm{high}}}$, and tracks $\boldsymbol{X}$ 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 $\boldsymbol{X}$.
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:

- $L_1(\boldsymbol{x}_t, \hat{\boldsymbol{x}}_t) = w_{\text{pos}}\gamma^t||\hat{\boldsymbol{\theta}}_{\text{S},t} - \boldsymbol{\theta}_{\text{S},t}||_2^2 + w_{\text{vel}}||\hat{\dot{\boldsymbol{\theta}}}_{\text{S},t} - \dot{\boldsymbol{\theta}}_{\text{S},t}||_2^2$, 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;
- $L_2(\hat{\boldsymbol{x}}_t, \boldsymbol{u}_t) = w_{\text{torq}}||\boldsymbol{u} - f_{\text{ID}(\hat{\boldsymbol{x}}_t,0)}||_2^2$, accounting for the difference between the generalized torques exerted on the human model and the torques that would guarantee equilibrium of the future trajectory;
- $L_3(\hat{\dot{\boldsymbol{x}}}_t) = w_{\text{acc}}||\hat{\dot{\boldsymbol{\theta}}}_{\text{S},t}||_2^2$, weighting the instantaneous acceleration of the human model.

Constraints: the set of constraints that we employ formalize the requirements that $\boldsymbol{X}$ 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 $\boldsymbol{\theta}_{\text{S},t} \notin \mathbb{Z}_U$;
- *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 $|\dot{\boldsymbol{\theta}}_{\text{S},T}| \leq \varepsilon_{\text{vel}}$, $||\boldsymbol{u}_T - f_{\text{ID}}(\boldsymbol{x}_T, 0)||_2^2 \leq \varepsilon_{\text{torq}}$, 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: $\text{dist}(\boldsymbol{x}_T, \mathbb{C}_U) \geq \delta$.

Through orthogonal collocation techniques [150], 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 equations (5.2) and (5.3) could introduce more involved computations without significantly affecting the quality of the optimal deflection $\boldsymbol{X}$. 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 5.3.1.

5.2.4. 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.

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**
  $\boldsymbol{K}_t \leftarrow \boldsymbol{K}_{high}$
  $\boldsymbol{D}_t \leftarrow 2\sqrt{\boldsymbol{K}_{high}}$
  $\boldsymbol{X}_{ref} \leftarrow \textbf{argmin}_{\boldsymbol{X}_{\text{ref}}} \mathscr{L}(\boldsymbol{X}, \boldsymbol{X}_{\text{ref}})$
else
  $\boldsymbol{K}_t \leftarrow 0$
  $\boldsymbol{D}_t \leftarrow 0$
end if
Return: $[K,\ D,\ \boldsymbol{p}_{\text{EE}}(\boldsymbol{X}_{\text{ref}})]$ to Robot Control

5.2.5. 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 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 5.2.2. 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 $k_{\text{low,t}} = 20\frac{N}{m}$ and $k_{\text{low,r}} = 5\frac{N}{[rad]}$.

By removing the damped zone in the low-authority module, we effectively recreated the haptic boundaries considered in our previous work [61], to measure the performances of this baseline controller in terms of robot-commanded forces when the unsafe zone was reached (Section 5.3.2). 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.

5.3. Results

5.3.1. Evaluation of the approximated dynamics

We employed computer simulations to evaluate the differences between the optimal deflections produced by the NLP presented in Section 5.2.3 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 5.3 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.

5.3.2. Baseline controller

We ran our baseline controller [61] 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.

5.3.3. 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's behavior was dictated alternatively by one of the two algorithms presented in Sections 5.2.3 and 5.2.3. 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 5.4a) 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 5.4b), the subject could move undisturbed closer to the unsafe zone until their predicted future trajectory was no longer safe. At this

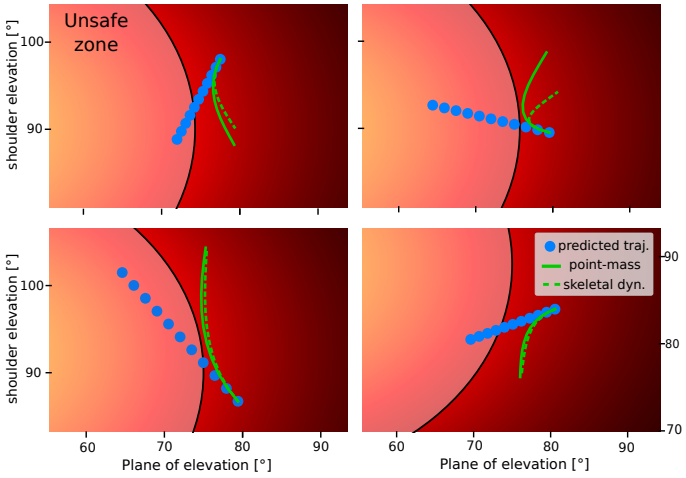

Figure 5.3: *Simulation comparison of optimal deflections generated by the high-authority module, contrasting the results when using the human model's dynamics (dashed green) versus a virtual point mass (solid green) for navigating the strain map.*

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.

5.4. 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 [60] and did not anticipate human intentions [61].

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 5.4a). 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

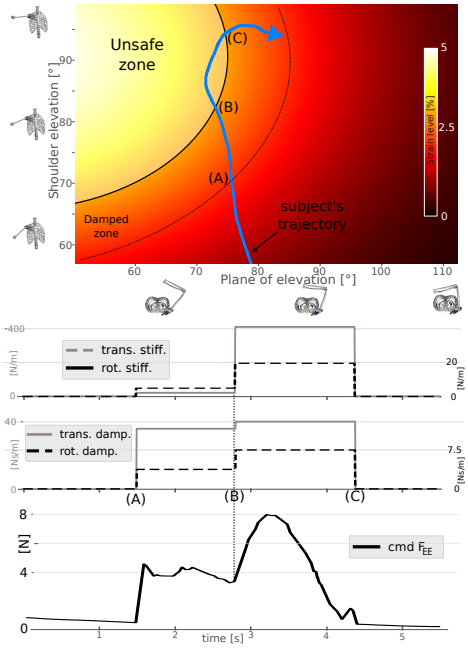

(a) *Low-authority shared control module behavior. Top: strain map in which the subject moved, with skeletal poses detailing the axes. The subject approached the unsafe zone: (A) first, the "damped zone" was encountered, then (B) they entered the "unsafe zone" and experienced higher robot correction 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.*

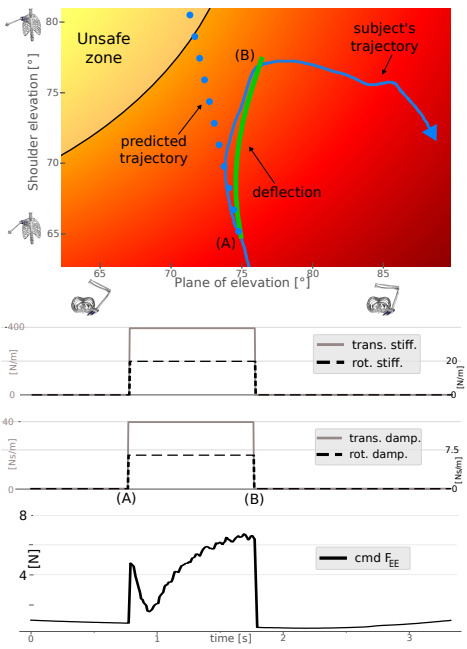

(b) *High-authority shared control module behavior. Top: strain map in which the subject moved, with skeletal poses detailing the axes. The subject approached the unsafe zone, triggering the system to impose a short deflection on their movement (A). After that (B), control was given back to the human,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

Figure 5.4: *Comparison of low- and high-authority shared control module behaviors during physical human-robot interaction.*

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 5.3, 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 5.3), likely due to different relative weighting of the torque terms in the cost function. Indeed, the inertial parameters of the virtual point mass navigating across the strain map were not tuned to our participant since only the human kinematics were 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 [60].

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 [147]. Including rapid estimation of the muscle activation through dedicated solvers such as [59] 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.

6

Recording patient's discomfort to augment computational biomechanics

Previous chapters have proposed methods to integrate musculoskeletal models and derived metrics into robotic control, enhancing physiotherapy across different stages. By leveraging these metrics, robots can precisely estimate how their interactions with patients affect the human musculoskeletal system, enabling safer and more effective treatment for complex injuries such as shoulder rotator cuff tears.

Yet, while these quantitative metrics help assess re-injury risk, they do not capture the patient's subjective experience of discomfort, or even pain, during therapy. Estimating discomfort reliably is difficult, given its variability across individuals, influenced by factors like inflammation and personal sensitivity.

To address this limitation, this chapter augments computational biomechanics, introducing an intuitive method to log patients' discomfort information. We enable patients to continuously communicate discomfort via a handheld push-button device, while interacting with a robotic physical therapy device capable of moving the patient and estimating their pose. To ensure safety and provide ground-truth feedback, we emulated discomfort via an auditory signal, allowing participants to reconstruct synthetic "discomfort maps" that demonstrate the feasibility of automatically capturing and storing patient perceptions during robotic physiotherapy.

Our discomfort maps have the potential to increase the feasibility of robotic therapy, while unlocking new diagnostic opportunities. When integrated with other metrics such as the strain maps, they can provide a more comprehensive representation of a patient's mobility while accounting for both objective biomechanical quantities and subjective perceptions during rehabilitation.

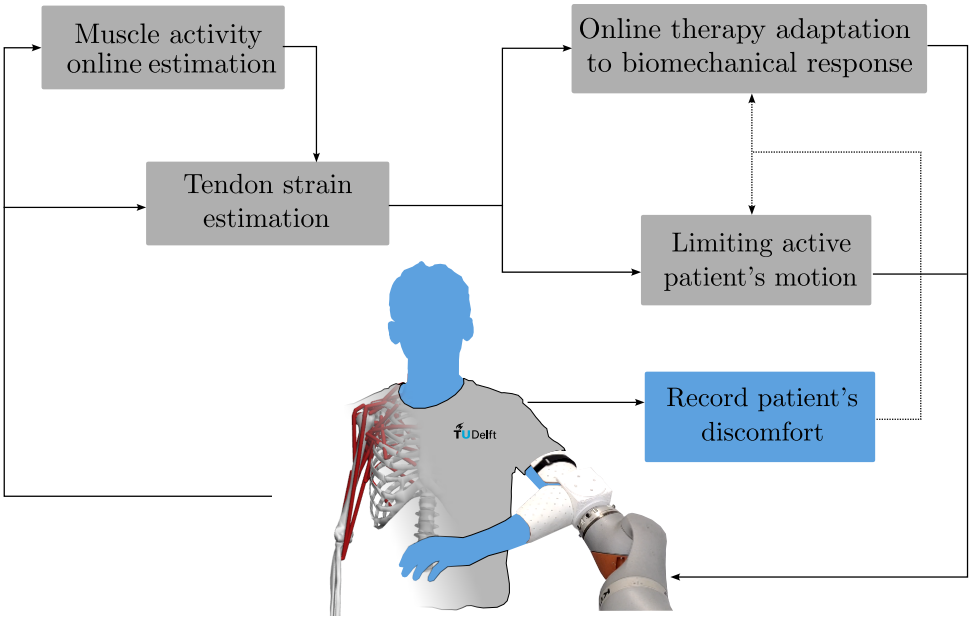

This chapter is based on J. Ravenberg, **I. Belli**, J.M. Prendergast, A. Seth, and L. Peternel. Creating Discomfort Maps via Hand-held Human Feedback Interface for Robotic Shoulder Physiotherapy. 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 4664-4671. [192].

6.1. Introduction

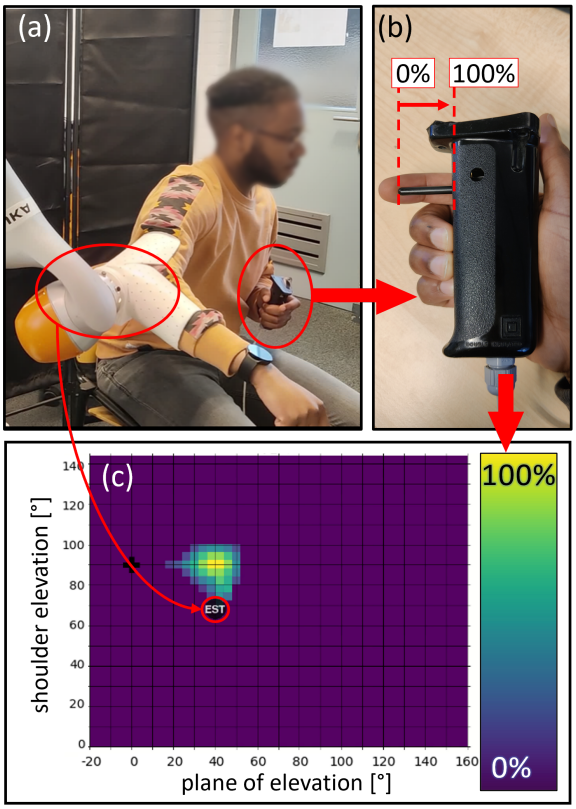

Figure 6.1: *Overview of the discomfort map creation system. (a) A KUKA LBR iiwa 7 R800 robotic manipulator acts as a physiotherapist and moves the participant's shoulder while simultaneously estimating their shoulder state. (b) A push-button interface (linear potentiometer), held by the free arm, is used to input discomfort intensity. (c) Example of a discomfort map created during the robotic shoulder physiotherapy.*

Musculoskeletal injuries resulting from accidents, recreational activities, and aging are the primary contributors to disability and work impairment. Among these injuries, shoulder rotator-cuff (RC) tears stand out as one of the most prevalent, with an estimated prevalence rate of 22.1% in the general population and over 50% for individuals aged 60 and above [193]. Restoring shoulder mobility and functionality after RC injuries requires a patient to undergo a prolonged and costly physiotherapy process. Due to the complexity of the shoulder mechanism and a lack of quantitative insights into the risks of re-injury, conventional practices in RC physiotherapy tend to be conservative, even when administered by expert physiotherapists [194]. This conservative approach limits the treatment range of motion (RoM), potentially prolonging the recovery process. Increasing RoM safely, however, can enhance recovery speed

and completeness [195]. Additionally, administering physical therapy can be physically demanding for physiotherapists who typically treat many patients a day, one at a time. Demand on physiotherapists is exacerbated by the growing gap between the number of people with rehabilitative needs and the number of available physiotherapists, due in part to the aging population and decreasing availability of medical personnel [196].

Robotic-assisted rehabilitation can offer solutions to both the safety and delivery challenges of rotator-cuff-related physical therapy. Robotics assists in reducing the physical load on the physiotherapist by taking over the weight of the patient's limbs, and potentially enabling therapists to operate remotely or with multiple patients at once. Furthermore, robotic platforms can quantify the patient's condition as input into their control system, and communicate vital information to the therapist. For instance, the muscle effort of the patient can be estimated from electromyography (EMG) measurements [197, 198], while musculoskeletal models can give an accurate estimate of the actual internal properties of the human body, such as joint loading [199], muscle fatigue [200], muscle comfort [201], and muscle manipulability [202]. Previously, our group has used a musculoskeletal model of the human shoulder to estimate safe RoM for rehabilitation, identifying configurations where the strain in the RC tendons are indicative of increased re-injury risk [203, 204]. The human RoM was abstracted into "strain maps", providing an intuitive representation of RC tendon strains and enabling more accurate monitoring of tendon strains by a physiotherapist and/or robotic system [147, 203, 204].

While model-based methods provide useful information about the patient's biomechanics, they do not quantify the patient's experience of discomfort or pain that might occur during a therapy session. Musculoskeletal models provide insights into mechanical variables like tendon strain, which are related to discomfort, but estimating discomfort itself remains challenging due to the wide variability among individuals, influenced by factors like inflammation and personal sensitivities . [205]. Despite being challenging, a patient's discomfort cannot be disregarded, and methods to personalize therapy to the experience of discomfort are necessary for robot-mediated physiotherapy to be acceptable.

Perceived pain can be recorded based on physiological markers such as blood pressure, heart rate, and skin conductance [206], or behavioral responses such as facial expressions, which may be useful for informing machine learning algorithms [207]. Cheaper and more portable interfaces can allow the patient to log their perceived pain level directly, for example using a slider [208, 209], or single-handed grip devices as the input interfaces [210, 211], which could allow a therapist to compare the discomfort that the patient feels against subjective numerical ratings, e.g., the Numerical Rating Scale (NRS) [212]. In the rehabilitative robotics realm, limited work has addressed the problem of endowing robots with awareness about the pain/discomfort that their patients are experiencing. Perceived pain levels have been monitored after the therapy through the NRS [213], and online approaches have been designed with very precise hardware in mind [214] or proved to be challenging to generalize, requiring careful validation and training of the proposed pain-detecting algorithms [215].

Drawing inspiration from the interaction of physiotherapists and their patients, we propose a new method of integrating patient responses to perceived discomfort during

robotic physiotherapy, to create "discomfort maps" that identify discomfort levels at different poses of the patient. Our main aim is to enable patients to autonomously provide their discomfort level in real-time during the robotic therapy session through a handheld push-button interface (see Figure 6.1), and create discomfort maps that represent the discomfort distribution experienced over the RoM of the therapy. This information is inherently personalized to the patient and offers quantitative data, for informing both robot and human therapists about their patients' perceptions in real-time. By using a linear potentiometer as the push-button interface, our method combines the single-handed, intuitive nature of the grip interfaces with the more precise position-based and patient-agnostic continuous input of the slider.
The main contributions of this study are threefold:

1. we present the concept of discomfort maps and introduce a novel combination of a handheld device with a collaborative robot to log patient discomfort during RC therapy as a function of their movements;

2. we validate our system through human factors experiments using auditory signals to emulate discomfort during movement, and measure user responses to test the accuracy of users in reconstructing a prescribed map;

3. we propose that the subjective perception of discomfort could be used to enhance and further personalize other existing tools for robotic-assisted rehabilitation, such as the strain maps, and enable tracking of patient progress while unlocking new diagnostic opportunities.

6.2. System Design

Our system is designed to allow the patient to input in real-time their level of perceived discomfort, as a continuous variable, using a push-button interface. Then, personalized discomfort maps are generated by relating the discomfort intensity to the current patient's shoulder state. To estimate the patient's shoulder state we employed a collaborative robot arm, which can also move the patient's shoulder through set trajectories (Figure 6.1).

6.2.1. Shoulder State Definition

In the context of RC tendon rehabilitation, we define the shoulder state as the state of the glenohumeral joint [203, 204], i.e., the motion of the humerus (upper arm bone) relative to the scapula (shoulder blade), which has 3 degrees of freedom (DoF). As such, the shoulder state vector is defined as

$$\boldsymbol{\theta}_S = [\text{AR PE SE}], \tag{6.1}$$

where AR $\in [-90°, 90°]$ is the axial rotation, PE $\in [-20°, 160°]$ is the shoulder plane of elevation, and SE $\in [0°, 144°]$ is the shoulder elevation. These coordinates are shown in Figure 6.2. The state $[0°, 0°, 0°]$ coincides with the neutral pose where the arm rests at the side. The continuous shoulder range is discretized in 4° increments.

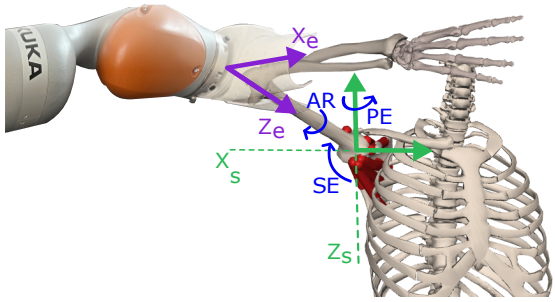

Figure 6.2: *Definitions of shoulder and end-effector coordinate reference frames. The shoulder frame (green), with the origin centered on the glenohumeral joint; DoF of the glenohumeral joint (blue); robot's end-effector frame (purple). When the patient is wearing the arm brace, the elbow frame coincides with the end-effector frame.*

6.2.2. Discomfort Map

Similar to our previous efforts using strain maps [203, 204], we develop the idea of discomfort maps to indicate the intensity of a patient's discomfort over the RoM of interest. The idea behind the map creation is that the patient indicates shoulder configurations that cause discomfort, and the intensity of this discomfort. Our simultaneous logging of patient's discomfort and estimation of the patient's current shoulder state allows us to map the continuous, variable discomfort intensities directly onto the patient's RoM, creating personalized discomfort maps. For visualization purposes, the 3D map of the patient's discomfort intensity level across AR, PE, and SE is represented as multiple layers of 2D maps, where AR is fixed. A single discomfort map shows the patient's discomfort intensity across PE and SE for a fixed AR, and by layering discomfort maps for all values of AR the entire shoulder RoM is spanned. An example of such maps is given in Figure 6.3.

When creating a discomfort map, we assume that the perceived discomfort should not be specific to the current shoulder state, but also affect neighboring poses. Upon patient input, 3D ellipsoidal Gaussian clouds are logged in the discomfort maps, scaling them in amplitude and size with the intensity input, and centering them around the current shoulder state.

As such, the effect of the discomfort intensity $i(t)$ that the patient inputs on the information that the discomfort map carries at a generic shoulder state $\boldsymbol{\theta}_S$ at time t is:

$$I(t, \boldsymbol{\theta}_S) = \max\left(I(t - t_s, \boldsymbol{\theta}_S),\ \ f(t, \boldsymbol{\theta}_S)\right), \tag{6.2}$$

$$f(t, \boldsymbol{\theta}_S) = i(t) \prod_{j=1}^{3} \exp\left(-\frac{1}{2}\left(\frac{\theta_{S,j} - \hat{\theta}_{S,j}(t)}{s_j(t)}\right)^2\right), \tag{6.3}$$

$$s(t) = i(t)\frac{1}{4}[w_{AR}\ w_{PE}\ w_{SE}] \tag{6.4}$$

where $I(t, \boldsymbol{\theta}_S)$ is the resulting discomfort level for shoulder state $\boldsymbol{\theta}_S$ captured in the map at time t, t_s is the sampling time, $f(t, \boldsymbol{\theta}_S)$ describes the intensity distribution of

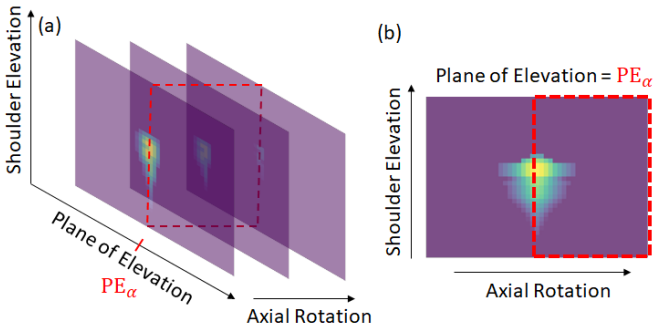

Figure 6.3: *Example of our discomfort maps: (a) maps for multiple values of AR layered on each other; (b) a cross-sectional view of the bounded area in (a), where the plane of elevation is at a fixed value of PE_θ. This demonstrates how the logged discomfort intensity dampens along the 3 DoF of the shoulder.*

the newest Gaussian cloud, $\hat{\boldsymbol{\theta}}_S(t)$ is the estimated patient shoulder state, and $s(t)$ collects the standard deviations of such Gaussian cloud along the 3 DoF of the shoulder, proportional to input intensity through $[w_{AR}\ w_{PE}\ w_{SE}] = [60°\ 25°\ 25°]$. We selected a larger Gaussian cloud width in the AR direction since a rotation along AR corresponds to a relatively smaller movement of the arm, compared to equivalent rotations in PE and SE. This adjustment aims at augmenting the safety bounds along the DoF along which unnoticeable movement is more likely to happen. Equation (6.2) describes how the logged discomfort intensity for each shoulder state $\boldsymbol{\theta}_S$ can only be overwritten by a Gaussian cloud with higher intensity. This stems from a conservative approach that regards the highest intensity input at the same shoulder state as the most representative.

6.2.3. Patient Input and Feedback Interfaces

The interface used for inputting discomfort intensity must enable the patient to instantly indicate not only whether and when they experience discomfort, but also to quantify the level of such discomfort. Furthermore, the interface should be comfortable to hold for a long time and intuitive. To comply with these requirements, we use a handle with a spring-return linear potentiometer, already employed effectively in other human-robot interaction tasks [216] (Figure 6.1.b). The discomfort intensity $i(t)$ input to the system scales linearly with the push-button position, ranging from 0% when fully extended to 100% when fully pushed in. The total allowable travel of the push-button is 2.5 cm, allowing for fine resolution in the discomfort logging. Further, participants held the handle in their free arm, to prevent the effects of involuntary muscle contractions due to the therapy on the button press.

6.2.4. Robot Control

A collaborative robot arm is used to estimate the shoulder state of the participant through the robot's built-in encoders, while also guiding the participant's shoulder through set trajectories. For estimating the human shoulder state, we follow the same approach as our previous work, as the shoulder angles are fully determined by the end-effector pose when the glenohumeral joint center does not move [203, 204]. During the experiments, we monitored our participants to satisfy this condition. To execute the experimental trajectories, an impedance controller is used, which takes the desired motion of the robot end-effector as an input and calculates the required robot joint torques based on mass-spring-damper equations [191]. Unlike fully stiff conventional position-controlled robots, the torque-controlled robot with an end-effector impedance controller can be compliant, which is crucial for safe human-robot interaction. The robot end-effector force and corresponding torques that generate that force are calculated with:

$$\mathbf{F} = \mathbf{K}(\bar{\mathbf{x}}_{EE} - \mathbf{x}_{EE}) - \mathbf{D}(\bar{\dot{\mathbf{x}}}_{EE} - \dot{\mathbf{x}}_{EE}) \tag{6.5}$$

$$\boldsymbol{\tau} = \mathbf{J}^T\mathbf{F} \tag{6.6}$$

where $\mathbf{F} \in \mathbb{R}^6$ is the output force, $\mathbf{K} \in \mathbb{R}^{6\times6}$ is the Cartesian stiffness matrix, $\mathbf{D} \in \mathbb{R}^{6\times6}$ is the Cartesian damping matrix, $\bar{\mathbf{x}}_{EE}, \bar{\dot{\mathbf{x}}}_{EE} \in \mathbb{R}^6$ are the Cartesian end-effector pose and velocity references, $\mathbf{x}_{EE}, \dot{\mathbf{x}}_{EE} \in \mathbb{R}^6$ are the current end-effector pose and velocities, $\mathbf{J} \in \mathbb{R}^{6\times7}$ is the robot Jacobian, which describes the relationship between end-effector velocities and joint velocities, and $\boldsymbol{\tau} \in \mathbb{R}^7$ is the commanded robot joint torques. The system is critically damped so $\mathbf{D} = 2\sqrt{\mathbf{K}}$ [217].
To ensure repeatability, a null space controller is also employed [218], which drives the robot towards a desired joint configuration. The null space torque is described by:

$$\boldsymbol{\tau}_{null} = (\mathbf{I} - \mathbf{J}^T\mathbf{J}^{+T})(\mathbf{P}_N(\bar{\boldsymbol{\theta}}_R - \boldsymbol{\theta}_R) - \mathbf{D}_N(\dot{\mathbf{q}})) \tag{6.7}$$

where $\mathbf{I} \in \mathbb{R}^{7\times7}$ is the identity matrix, $\mathbf{J}^+ \in \mathbb{R}^{7\times6}$ is the Moore-Penrose pseudo-inverse of the Jacobian $\mathbf{J}$, $\mathbf{P}_N \in \mathbb{R}^{7\times7}$ is the proportional gain of the controller, $\mathbf{D}_N \in \mathbb{R}^{7\times7}$ is the derivative gain (damping), $\bar{\boldsymbol{\theta}}_R \in \mathbb{R}^7$ is the reference joint configuration, and $\boldsymbol{\theta}_R \in \mathbb{R}^7$ is the current joint configuration of the robot. Also here, we selected a critical damping $\mathbf{D}_N = 2\sqrt{\mathbf{P}_N}$. The torque commanded to the low-level controller is thus $\boldsymbol{\tau} + \boldsymbol{\tau}_{null}$.

6.3. Experiments and Results

We separate our experiments into two main parts:

- functionality test (Section 6.3.3), with the goal of demonstrating the discomfort map creation process;
- human factors experiments (Section 6.3.4), with the goal of analyzing the practicality of the proposed method on untrained participants.

Ten healthy participants (8 male and 2 female) in the age range of 20-30 participated in our study. Our experimental protocol was approved by the TU Delft Human Research

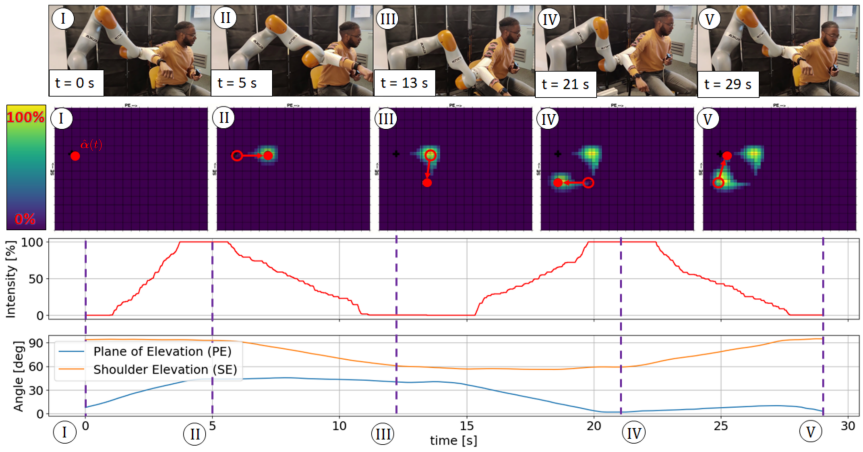

Figure 6.4: *Functionality test demonstrating a discomfort map creation, with snapshots taken at specific timestamps. The top row shows the robot and participant poses. The second row shows the gradual creation of the discomfort map, with the current shoulder state marked as a filled red dot and the previous as a non-filled circle. The third row shows the discomfort intensity input* $i(t)$ *from the participant over time, i.e., how much the push-button was pressed. The bottom row shows the PE and SE angles over time. AR is missing from this overview because, for this demonstration, the estimated AR was locked to 0. Note: The images in the first row were mirrored for demonstration purposes.*

Ethics Committee (HREC) and all participants provided informed consent before their participation. For HREC approval, we could not induce (physical) discomfort on the participants to generate discomfort maps in the human factors experiments. Therefore, we emulated discomfort with an auditory feedback signal provided to the participant. This signal was chosen to be a beeping sound pattern, further explained in Section 6.3.2. We set *a priori* an artificial reference discomfort distribution across the range of motion of our participants and asked them to recreate this reference using the proposed method in a simulated physiotherapy session, during which their shoulder was moved by the robot arm. They were not able to see the reference map, but were instead given a reference auditory signal based on this map, emulating the corresponding discomfort and dependent on their current configuration. In this way, we intended to reproduce a situation in which all of our participants would perceive the same discomfort when performing the same motion, to assess how well a patient could synthesize an actual discomfort map in the context of a real rehabilitation (with the assumption that a discomfort distribution for this case would still exist, yet unknown for both the patient and our system). As such, we instructed our participants to react to variations in the sound pattern as they would react to variations in their own physical discomfort.

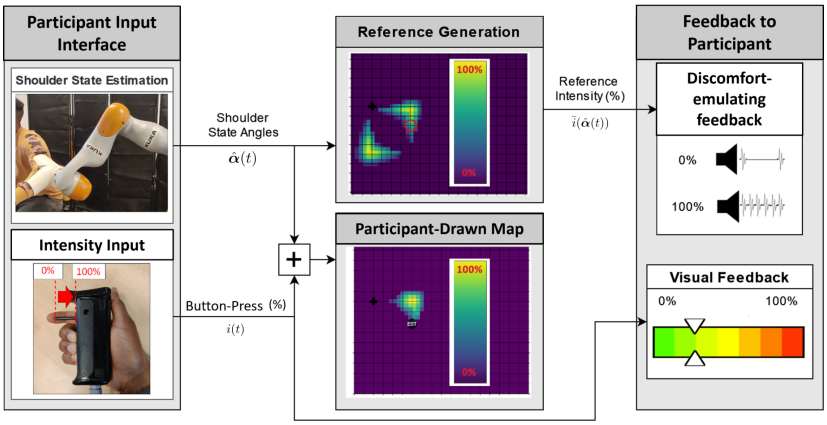

Figure 6.5: *Workflow for human factors experiments: Participants recreate a reference discomfort map guided by the robot and discomfort-emulating auditory feedback (Section 6.3.2). The system includes participant input, reference generation, map synthesis, and feedback. Participant inputs are their shoulder state (estimated by the robot) and perceived discomfort intensity (button press). The reference discomfort intensity, derived from the reference map at the current shoulder state, modulates an audio signal, with the beeping rate increasing with intensity. Feedback includes this audio signal and visual feedback of the participant's input. When drawing the discomfort map, the user generates Gaussian clouds centered at their current estimated shoulder state, whose shape depends on the participant's button-press input. The drawn discomfort map is not shown to the participant, to minimize distractions during the task, and is shown here as it appears in the process of drawing it.*

6.3.1. Experimental setup

During the experiments, participants were seated in a chair, adjustable in height, with their left elbow fitted in the elbow brace attached to the end-effector of the robot arm. The starting pose for every experiment was [AR PE SE] = [0° 0° 90°], where the bent arm is parallel to the horizontal plane (see Figure 6.4.I). The robot arm used is a KUKA LBR iiwa 7 R800 collaborative robot arm. The participants received visual feedback on their estimated shoulder state, the robot's reference trajectory, and their input discomfort intensity through a monitor. The participant also received the auditory signal through Bluetooth headphones in the second part of the experiment.

For all the experiments, AR was locked to 0°, reducing the complexity of the motions to be performed. This means that the discomfort maps coincide with the discomfort distribution considered, improving the visualization of the results.

6.3.2. Discomfort-emulating Auditory Signal Protocol

As stated before, during the human factors experiments each participant received a signal emulating the reference discomfort intensity $\bar{i}(\hat{\boldsymbol{\theta}}_S(t))$, retrieved from the artificial

reference discomfort distribution given their current estimated shoulder pose. We chose sound feedback over visual feedback so that the participant could focus visual attention on the scene and human-robot interaction. We chose sound feedback over tactile feedback as it can be delivered in a more repeatable and clear manner. Furthermore, tactile human-robot interaction might interfere with or confuse tactile feedback.
To convey changes in reference intensity, we considered two types of auditory signals: continuous sound or beeping sound, where the former can vary in volume/frequency, while the latter can vary the rate of the beeping. After preliminary testing, volume/frequency modulation of a continuous sound was found to be more responsive but much less clear. This may be because humans do not perceive loudness linearly, and beeping rate modulation is easier to recognize than frequency modulation. Thus, finally, we chose the beeping sound for its clarity advantage at the expense of reaction time. The auditory signal was implemented as a repeating 1 kHz beep sound with modulation of beeping rate. The time τ_b between beeps, expressed in seconds, was inversely proportional to $\bar{i}(\hat{\boldsymbol{\theta}}_S(t))$, i.e. the emulated discomfort at the current estimated shoulder state $\hat{\boldsymbol{\theta}}_S$, and varies according to the law:

$$\tau_b = (1 - \bar{i}(\hat{\boldsymbol{\theta}}_S(t))) + \tau_0 \tag{6.8}$$

where $\tau_0 = 0.3$ seconds is the duration of the beep itself, constant and independent of the reference.

6.3.3. Functionality Test

The goal of the functionality test was to demonstrate the discomfort map creation process in the space of the shoulder DoFs. During this task, the robot was fully compliant, i.e., stiffness and damping parameters for the impedance control were all set to 0, and its only role was to estimate the shoulder state during the experiment. Figure 6.4 shows how a discomfort map could be created autonomously by a subject during a simple motion of the left arm.

6.3.4. Human Factors Experiments

The goal of the human factors experiments was to analyze the practicality of the proposed method on untrained participants. Figure 6.5 gives a general overview of the system. The experiment simulated a robotic physiotherapy session, in which the robot arm guided the participant's shoulder through a reference trajectory, while the participant was asked to passively follow the movement commanded by the robot. During the fictitious therapy, the pose of the robot's end-effector was used to estimate the current shoulder state $\hat{\boldsymbol{\theta}}_S$, and the beep-rate of the discomfort-emulating sound was then retrieved according to equation (6.8). Note that the combination of impedance control and physical interaction with the subject caused the robot to follow slightly different trajectories for each experiment execution, meaning that the subjects received different discomfort-emulating auditory references even when the robot's controller tracked the same trajectory more than once. The reference maps and corresponding trajectories

(henceforth collectively referred to as "reference") used for the tasks are shown in Figure 6.6, labeled A, B, and C.
The experiments began with a familiarization period, where the participant could freely press the button to explore how its linear displacement corresponded to the color bar for visual feedback and the beeping rate of the discomfort-emulating signal. This step was crucial for the main experiment, where participants needed to match their button press to the signal. After taking as much time as needed, they wore the arm brace and were connected to the robot. Each participant repeated every task 5 times, starting with the easiest (task A) and progressing to the most complex (task C). For safety, the robot's controller gains were tuned such that the robot could move the participant's arm while ensuring that the participant would still easily overpower it. Task repetitions were limited to prevent memorization of reference signals. Breaks were allowed, with each experiment lasting about 20 minutes.
To get quantitative insights into the similarity between the artificial reference discomfort distributions and the maps that the participants reconstructed, we employed two metrics:

- the intensity score, describing how the participant's input intensity $i(t)$ matches the reference intensity instantaneously;
- the map score, describing how the final discomfort map drawn by the participant matches the reference map.

We computed both scores as the root mean square error (RMSE) between the reference and the participants' responses. The intensity score essentially represents the average deviation between the instantaneous reference intensity $\bar{i}(\hat{\boldsymbol{\theta}}_{\mathrm{S}}(t))$ and $i(t)$ over the duration of the experiment. On the other hand, the map scores account for the point-wise RMSE between the map that participants reconstructed and the reference map, evaluated on regions of the maps where either was non-zero.

6.3.5. Results

Concerning the functionality test, Figure 6.4 demonstrates that the combination of our hand-held device with a human-pose estimation from the collaborative robot enables the participant to transfer their perceived discomfort into a quantitative measure that is directly mapped onto the RoM of their shoulder. For what concerns the outcomes of the human factor experiment, Figure 6.6 shows qualitatively the resulting discomfort maps, obtained as an average across the various participants, for each of the given trajectories. As the discomfort-emulating auditory signal was dependent on the participant's current position on the reference map, the reference they received was subject to some variability. This is visible in Figure 6.7, where we present the distribution of references and participants' responses (button press) against the normalized duration of the experiments. In the same figure, it is evident that the participants reacted with some delay to the discomfort-emulating feedback. On average, this was observed to be around 1.4 seconds.
The average scores are shown in Table 6.1, where lower scores indicate better performances. Performance for reference A is the best out of all references, which makes sense as it is the simplest trajectory. As expected, the presence of an error in

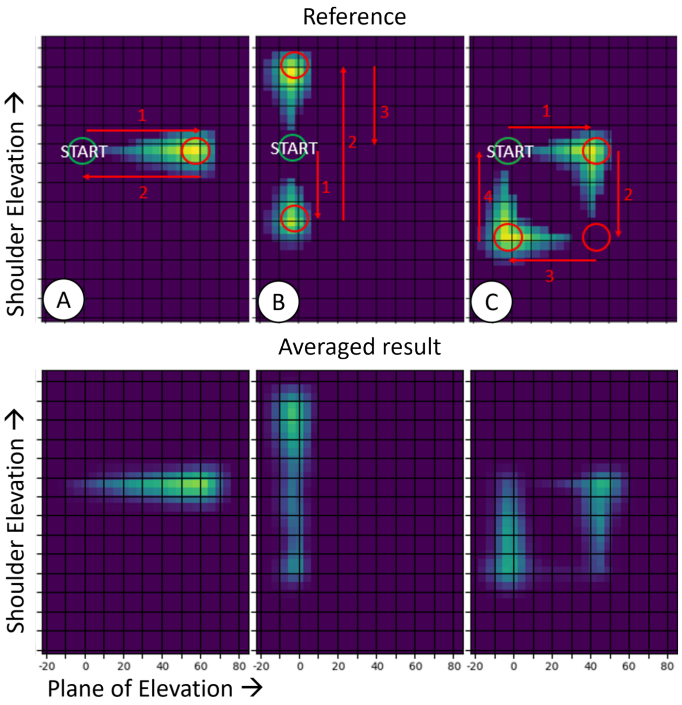

Figure 6.6: *The average discomfort maps created by the participants on the top row and the reference maps on the bottom row. The trajectories (A,B and C) were executed by the robot, and are here visualized in the shoulder space with red arrows, where the numbers clarify the sequence of execution for each phase of the movement.*

the tracking of the instantaneous discomfort-emulating sound directly translates into a poorer performance in the map scores. The low standard deviations reported in the table highlight that, on average, all the participants performed quite similarly across the different tasks, and the multiple repetitions for each task.

6.4. Discussion

The human factor experiments identified important insights for the practical use of the proposed method. We can see in Figure 6.6 that the averaged maps resulting from the participants' inputs do not match the reference completely (especially for tasks B and C). This is also evident when looking at the temporal relationship between the participants' input through the hand-held device and the reference discomfort-simulating feedback, in Figure 6.7. It can be seen that the participants consistently underestimated the percentage of button presses that were expected from them, while also introducing a significant delay. We cannot speculate if similar values would be observed when participants experience real discomfort, but we can hypothesize that these discrepancies

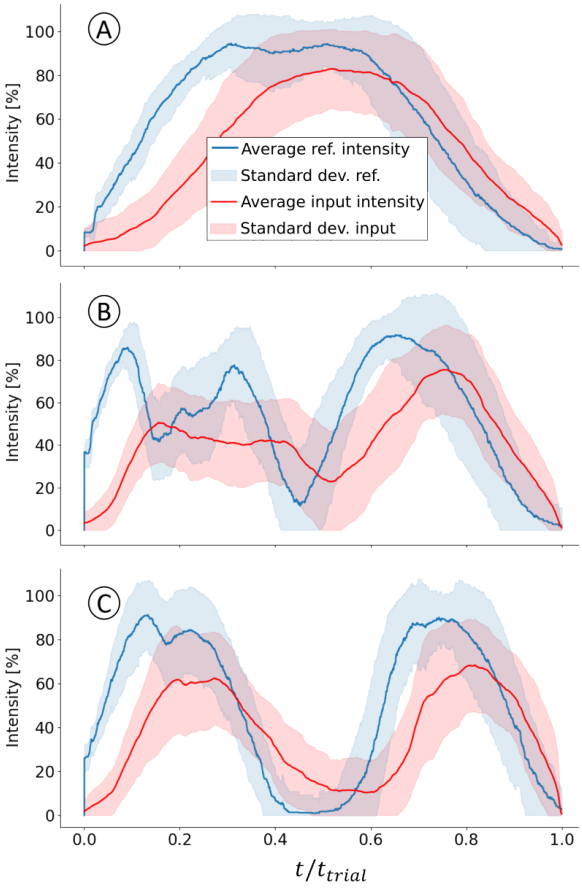

Figure 6.7: *Plots of the average discomfort intensity input from the participants and the average reference intensity over normalized time. It can be observed that the reference intensity is different for every trial, as the participants' shoulder movement is not precisely the same for each repetition of the task. This also entails that every trial has a different duration, and as such the data is normalized with the duration of the trial, then averaged.*

Table 6.1: *Average intensity and map scores, indicating the RMSE between references and participants' input*

	Intensity Score			Map Score		
Task	A	B	C	A	B	C
Average	0.24	0.33	0.32	0.26	0.34	0.36
Std. dev.	0.05	0.03	0.05	0.06	0.03	0.03

could be caused by the auditory feedback as a medium to deliver the reference for the emulated discomfort. Since causing real discomfort was not possible within ethical constraints and discomfort had to be emulated, this remains a limitation of the study.

Our choice of employing a beep-rate modulation for the audio signal entails that, to detect a change, one should wait for at least 1 full beeping cycle (introducing an intrinsic delay that is dependent on the beeping rate itself). This delay, for low discomfort-emulating references, would result in $\tau_0 + \tau_b|_{\bar{i}=0} = 1.6$ seconds, which is close to the average delay of 1.4 seconds that we observed across our experiments. This delay should be effectively identified and eliminated in our maps before they can be used for real physiotherapy.

Regarding $i(t)$ consistently being lower than $\bar{i}(\hat{\boldsymbol{\theta}}_S(t))$, this could be because participants tended to press the button cautiously or due to the aforementioned delay. We believe that this would not invalidate the use of our system when capturing actual discomfort, as the button press would essentially be equivalent to continuous logging of a pain-measuring scale such as the NRS, already clinically accepted for pain quantification.

There are several key benefits to the presented approach. It enables patients to quickly and accurately convey points of discomfort and their intensity to the physiotherapist, surpassing the speed of verbal communication and the precision of body language. Furthermore, the push-button interface for discomfort input presented here improves upon similar devices from the studies in [208–211] by being single-handed, not requiring personalization, and directly integrated into a robotic system. Moreover, when integrated with a patient movement tracking system, it becomes possible to map their discomfort distribution based on the patient's pose. Though this study only addresses shoulder physiotherapy with a collaborative robot arm, this method could easily be adapted to other parts of the body or different types of robots, like exoskeletons.

In the context of robotic-assisted shoulder rehabilitation, the discomfort maps have the potential to be integrated with the strain maps [147, 203, 204], and make it possible to build a comprehensive representation of the mobility range of a patient while accounting for both objective biomechanical quantities and the patient's more subjective perceptions during rehabilitation. While strain maps are already personalized to the patient and the severity of their injury, discomfort maps allow further personalization that even accurate biomechanical models cannot provide to date. As an example, the discomfort maps can also consider sources of discomfort unrelated to RC tendon strain, such as inflammation in the shoulder, or different levels of pain acceptance for various patients. Furthermore, the simplicity of our pipeline enables us to easily update the discomfort maps as the therapy progresses, offering a quantitative tool that can inform the therapist, the robot-assisting device, and even the patients regarding the course of the recovery process.

Our method is also subject to limitations, mainly because we were not able to test its effectiveness in reconstructing discomfort maps that are directly related to patients' real perceived discomfort. While evaluating the quality of the discomfort maps in this real case would be challenging (as we would not have any ground truth information), future work should focus on evaluating the system as closely as possible to its real intended use. Another complication that might arise is that discomfort is likely

not dependent on the shoulder state $\boldsymbol{\theta}_S$ alone, but also on the speed at which the movements are performed, and on past movements as well. A system that is usable with real patients should also take into account these hystereses that are intrinsic to the physiology of the human body, and that are still an area of active research even in the medical community. Viable options to consider may be collecting discomfort information at various speeds and reaching painful poses from different directions, to identify how discomfort perception might vary and contribute to advancing our fundamental understanding of the human body.

6.5. Conclusions

We developed an intuitive interface to map patient discomfort during robot-assisted shoulder physiotherapy, enabling real-time quantification of discomfort as the shoulder is manipulated by a robotic device. This system allows us to track and mitigate discomfort directly, making robotic therapies more feasible. Moreover, the discomfort maps generated provide valuable documentation of therapy progress, aiding clinicians in tailoring rehabilitation programs. Discomfort information could be integrated with our previously developed strain maps, providing a comprehensive representation of the patient's mobility range in terms of both objective and subjective measurements. Analyzing the evolution of these maps throughout the rehabilitation process offers diagnostic insights and could inform treatment selection. While our system shows promise for broader application in robotic-assisted rehabilitation of various joints, further validation in clinical settings with real discomfort is necessary to fully assess its utility.

7

Discussion and conclusions

This final chapter synthesizes the contributions of this dissertation, reflecting on how musculoskeletal models and derived metrics can make rehabilitation robots more aware of the human body, thereby making them safer and more effective in physiotherapy. It offers insights into the proposed framework, in terms of both high- and low-level robot control, and how the developed methods could be combined in the future. We analyze current impact and identify future opportunities, from leveraging faster simulations and machine learning approximations to extending kinematic estimation and refining human–robot interaction strategies. Finally, we reflect on the broader outlook for physiotherapy: robots could complement rather than replace therapists, acting as biomechanically aware assistants that enhance safety, efficiency, and personalization of care, while opening new perspectives for human-centered technologies well beyond rehabilitation.

The work presented in this thesis advanced robotic-assisted physiotherapy with the design and development of control methods informed by simulations of the human musculoskeletal system. The main question that motivated my research was:

> *"How can musculoskeletal model–derived metrics be integrated into rehabilitation robot control online, shaping robot-mediated movement to prevent potentially harmful loading of rotator cuff tendons?"*

This question led to the development of a coherent approach to making robots aware of the inner functioning of the human shoulder, enabling automated patient monitoring and manipulation in real-time. We developed and tested several methods to improve the physical human-robot interaction (pHRI) between a rehabilitation robot and its patients, grounded in our novel online estimation of muscle activations and tendon strain in the human shoulder. We showed that this innovative combination of robotics and biomechanics enables improving the safety and efficacy of shoulder physiotherapy for patients with different needs: robots can assist early mobilization shortly after injury or surgery (Chapter 3), but also be helpful tools to monitor and protect autonomous movements (Chapters 4-5). In particular, we proposed a system that delivers patient-centered therapy, implementing control algorithms that make our robot reactive to the patient's biomechanical response (Chapter 3). Our contributions can have a significant impact on the state of the art for robotic rehabilitation, where robots were previously unaware of the underlying patient biomechanics [37, 128, 129, 219] or could not infer its changes during therapy execution [140].

7.1. Online estimation of musculoskeletal metrics

In this thesis, we contend that understanding precisely how the human body reacts to therapy is the key to achieving successful treatment for musculoskeletal injuries. While measuring certain quantities, such as range of motion or overall strength, can be fairly easy for a physiotherapist (PT), they lack access to deeper musculoskeletal insights.
Superficial human muscles and their recruitment during physiotherapy can be monitored through superficial electromyography (sEMG) [220, 221], and have already shown great success in ergonomics and robotics [222–225]. However, a key finding of our work was that the activity of deep muscles cannot be inferred from superficial signals alone (see Section 2.3.2 of this thesis). This limitation underscores the value of biomechanical-model-based approaches, which enable the estimation of deep muscle activations and related metrics that remain inaccessible to traditional sensing methods. Leveraging models and robotic sensing can allow us to design better rehabilitation protocols, and is pivotal to human and robotic understanding of how rehabilitation movements affect the underlying musculoskeletal system. In Chapter 4, for example, we showed that varying the level of robotic resistance during human-led shoulder training affected rotator cuff strains, with higher resistance actually decreasing strain for some exercises.
Our findings highlight how model-based estimates can support more informed clinical decisions for both patients and physiotherapists, while also opening avenues beyond

the rehabilitation context—such as ergonomic monitoring, sports performance analysis, and assistive robotics. Crucially, however, the full potential of these metrics emerges when they are not only observed but actively embedded into the operation of machines that physically interact with humans.

7.2. Musculoskeletal metrics in robot control

Deriving musculoskeletal metrics from models and sensor data online enables scientists and therapists to better understand the human body. An even more exciting opportunity lies in embedding this knowledge into machines that can regulate and assist the way humans move.

In the context of shoulder physiotherapy, this thesis leveraged a detailed biomechanical model to enable a collaborative robot to deliver safe robotic-assisted movements. In Chapter 3, we coupled musculotendon metrics into a receding horizon controller, introducing a novel biomechanics-aware trajectory optimization approach for online navigation of personalized strain maps (BATON).

Directly embedding a patient's biomechanical model in our trajectory optimization algorithm enabled the synthesis of movements tailored to the patient's tissue behavior, range of motion, and body inertia. As therapy progresses, models can be updated to reflect recovery, while control parameters can adapt to new therapeutic needs. This personalization also enables the robot to apply precisely the forces required for gravity compensation, safe accelerations, or strength training.

We showed that integrating accurate biomechanical models into a closed-loop controller is crucial for reacting to unforeseen human actions during therapy. Unlike previous work where biomechanical insights were restricted to offline planning [140, 144], our framework allows therapy to adapt in real time to fatigue, discomfort, non-optimal execution, or even sudden disturbances. Continuous feedback from online musculoskeletal metrics is therefore essential to ensure safety in the inherently unpredictable context of pHRI. Yet, it is important to recognize that musculoskeletal models can only capture objective biomechanical quantities, not the patient's subjective experience. Because discomfort and pain are deeply individualized, incorporating patients' self-reported perceptions —such as through our proposed *discomfort maps* (Chapter 6)— adds a crucial layer of personalization and is fundamental for the broader acceptance of robotic therapy. Surprisingly, this dimension has received little attention in rehabilitation robotics, despite its potential to guide control decisions alongside model-based metrics.

Taking a broader perspective, combining predictive biomechanical models with rapid closed-loop adaptation could transform assistive robotics altogether. In rehabilitation, we have demonstrated that robots could accommodate each patient's unique condition and range of motion, unlocking the synthesis of movements that protect tissues at risk. During early recovery from musculoskeletal injuries, robots could contribute to maintaining the range of motion and preventing muscle atrophy, while at later stages they could deliver tailored strength training. Neurorehabilitation could also benefit from the ability to isolate specific muscle groups, potentially enhancing motor learning for

people suffering from neurological conditions.
Beyond clinical settings, exoskeletons and prostheses could also profit from biomechanical-model-based control. In the workplace, passive and powered exoskeletons could ergonomically support movement by offloading fatigued structures without straining others. Such advances may help overcome the limited evidence currently supporting powered exoskeletons [226, 227] and contribute to preventing work-related musculoskeletal injuries. Similarly, active prostheses could leverage musculoskeletal models to better coordinate with the user's body, reducing parasitic forces and adapting mechanical properties to mimic those of healthy limbs.
In sum, embedding musculoskeletal awareness into control architectures opens the way toward machines that can actively understand and safeguard the human body they interact with.

Computational performances

Embedding musculoskeletal awareness into robot control is only achievable if relevant computations can be performed fast enough to keep up with human interaction. Indeed, at the beginning of this thesis, we found high computational load as one of the reasons for the limited applications of biomechanical models in robotics.
We addressed this limitation through a combination of novel estimation algorithms (RMR solver, Chapter 2), precomputed metrics (activation-dependent strain maps, Chapter 4), and the integration of recent software tools such as OpenSimAD [148]. A general strategy that proved effective in our work is to shift part of the computational burden offline: either by pre-computing quantities that can later be queried in real time, or by approximating complex physiological relations as functions of the changing model state. For example, we employed strain maps to encode the dependence of tendon strain on shoulder pose and muscle activations, and constructed their differentiable approximations to accelerate optimization algorithms. Other recent studies confirm the value of this approach, showing that similar approximations can accelerate predictions while preserving accuracy [32, 228]. Building on these ideas, we enabled our controller BATON to adapt reactively to patient (re)actions at about 10 Hz—already faster than the average human reaction time to somatosensory feedback, reported to be around 200 ms [229]. Our results show that musculoskeletal models can be employed online not just for analysis, but as active components of closed-loop control.
Nonetheless, scaling to more dynamic movements or to additional DoFs (such as scapula motion) will demand faster access to model outputs. One direction is the use of physics engines explicitly designed for performance, such as MyoSuite [230], though the scarcity of validated models currently limits their clinical relevance. Another avenue is the simplification of dynamics: in Chapter 5 we showed that a point-mass approximation can yield update rates of about 80 Hz, though such choices may not generalize to more complex cases. Thus, tighter integration of algorithms such as our RMR solver within OpenSim's C++ codebase could directly improve the performance of widely adopted, validated models. Additionally, parallel computations and adoption of gradient-free methods (like model predictive path integral control) [231, 232] could also help us cope with the computational demands of simulating accurate biomechanics.
Finally, it is worth noting that all our developments rely on open-source software tools

such as CasADi [149] and IPOPT [151], ensuring reproducibility and transparency. Commercial solvers (e.g., FORCESPRO [233]) could likely further improve computation times, albeit at the expense of openness.
In short, our results demonstrate that biomechanics-aware control is already feasible in real time for rehabilitation tasks, and that further improvements in computational performance will directly broaden the applicability of such approaches to more dynamic, higher-dimensional, and clinically realistic scenarios.

7.2.1. Robot control strategies

Throughout this thesis, our high-level controllers exploited efficient computations and model's abstractions to generate optimal robot references in real time. Specifically, BATON (Chapter 3) aimed at driving the human through safe movements defined by a desired final configuration for the human shoulder, which could represent either the end of the movement or an intermediate waypoint in a longer rehabilitation exercise. Chapter 5, instead, focused on keeping autonomous human shoulder movements away from unsafe configurations. This high-level collaborative goal was achieved either by synthesizing a safe alternative trajectory if required (without the need for input waypoints), or by employing haptic feedback to gently nudge the human away from unsafe configurations. The resulting output of our high-level controllers consisted of one or more desired human trajectories, that were then transformed into robot's end effector poses, which we tracked at the low level through our (variable) impedance controllers to allow for inherent compliance toward the human.
Our hierarchical control architecture mirrors common practices in robotic rehabilitation [20, 176], as it allows to decouple high-level, patient-specific therapy needs from robot-specific, low-level control. Indeed, our work focused mostly on the high-level layer, where musculoskeletal metrics can be most effectively embedded to shape therapy according to patient biomechanics, and leveraged impedance control [234] to ensure safe physical interaction at the low level. While impedance control is the most common paradigm nowadays for upper limb rehabilitation robots [235], its dual, namely admittance control, can also find meaningful applications to regulate physical human-robot interactions [236]. For example, outputting force as a consequence of detected human movement has great potential for assisting and augmenting patient rehabilitation, and has recently been employed for muscle-targeted robotic assistance [237]. Other interesting avenues arise from more complex selection and tuning of both impedance and admittance controllers' parameters, for example by assisting human movement initiation and termination by intervening dynamically on the damping's sign, in addition to regulating its magnitude [238].
The framework we presented is also general enough to accommodate alternative high-level control strategies, obtained not necessarily through optimal control. For example, human-robot trajectories could come from tele-operation by a remote therapist (such as implemented in our previous work [61]), and such trajectories could be combined with musculoskeletal simulations to generate novel rehabilitation movements through learning-based policies. Our safe trajectory deflection presented in Chapter 5 would be an ideal candidate to filter the output of such learning-based policies, preventing unsafe movements and implementing the biomechanic-centered equivalent

of predictive safety filters, which recently gained popularity in autonomous driving [239]. On the other hand, while our work focused predominantly on robotic assistance, control strategies inspired to BATON could also plan optimal robotic resistance, implementing challenge-based therapy [20] that targets specific muscle groups while keeping tendon strain within safe limits.
Finally, it is worth noting that real-time musculoskeletal metrics can also enhance more agile, portable cable-driven exoskeletons [46], opening the way for yet more advances in controllers for the future of physical human-robot interaction.

Models and AI

In this thesis, biomechanical models have been pivotal in enhancing robots' understanding of the human body. Yet, with the rapid rise of machine learning and artificial intelligence, both robotics and biomechanics are increasingly shifting from model-based approaches toward data-driven ones [240–242]. Properly trained function approximators, such as neural networks or Gaussian processes, can, in principle, provide biomechanical estimates more rapidly and directly from measurements, potentially surpassing physics-based musculoskeletal models in computational efficiency. For certain applications, this speed advantage may prove transformative. Additionally, some function approximators could provide a measure of the uncertainty bounds around their predictions, enabling the end user (human decision-maker or robotic controller) to weigh predictions against their associated confidence and adjust decisions accordingly. However, it is essential to recognize that machine learning models, in biomechanics such as in any other field, can only be trusted if trained on meaningful datasets and applied under conditions that remain sufficiently close to those represented in training. In this respect, accurate musculoskeletal models will remain indispensable. They provide the necessary foundation to inform, constrain, and validate machine learning models, ensuring their predictions are physiologically plausible and clinically safe. Looking ahead, such models may be employed offline to generate the large, diverse datasets required for effective machine learning training and for extensive sensitivity analysis, while continuing to supply critical patient-specific insights and otherwise inaccessible data—such as deep tendon strain during movement.

7.3. Future work

The work presented in this thesis contributed toward the vision of biomechanics-aware robotics, demonstrating how musculoskeletal models can be integrated into robot control and patient monitoring. Importantly, these considerations did not remain purely theoretical. A leading orthopedic physiotherapist in the Netherlands tested our system and highlighted its potential to prevent inadvertent tissue overload during therapy, as well as to serve as a diagnostic tool (Figure 7.1). Their endorsement provides clinical validation that biomechanics-aware robots could address pressing rehabilitation needs and may eventually find a natural place in practice.
Overall, our results point to several promising avenues for future work. In particular, I see three areas where opportunities are especially compelling: execution of user studies, improved model validation, and the use of robotics to advance biomechanics itself.

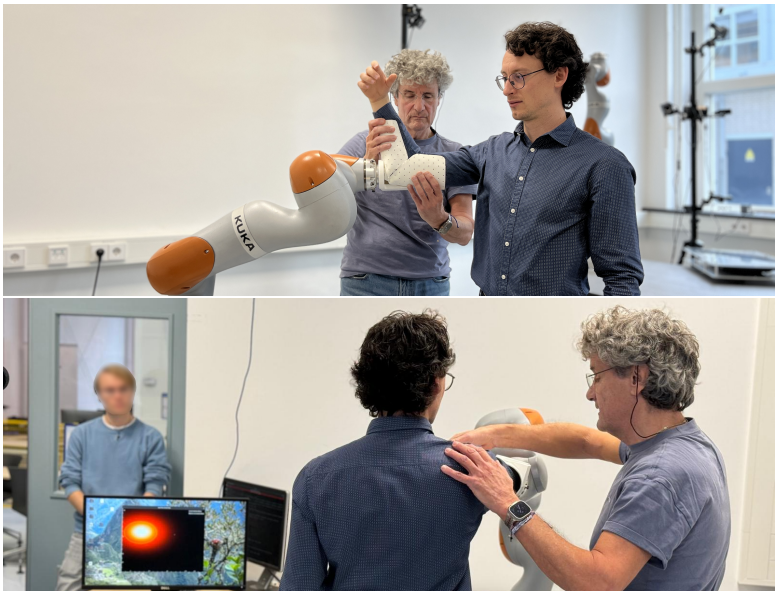

Figure 7.1: *A leading Dutch expert in shoulder physiotherapy administering therapy to a healthy subject in synergy with our system. Free movement was allowed when biomechanically safe, and reflected otherwise (with an early version of the work presented in Chapter 5). On the bottom left, our system can visually inform the physiotherapist about where tendon strain is unsafe.*

User studies

The methods developed here have not yet been validated through multi-participant user studies, which will be an essential step before extending our approach to clinical studies. User studies would not only test safety and efficacy at scale, but also allow fine-tuning of control algorithms based on the preferences of both patients and physiotherapists (PTs). Importantly, more research is needed on how patients perceive robotic guidance, particularly when the robot intervenes to correct unsafe movements, to ensure that assistance is perceived as supportive rather than conflicting. Understanding how to balance robotic authority with human autonomy will be key to improving both acceptance and therapeutic effectiveness. Moreover, cohort studies and assessment of clinical effectiveness are fundamental to stimulate further research in these topics, and eventually convince investors and decision makers of the promises that lie in human-centered rehab robotics.

Model validation and generalization

Another major opportunity lies in advancing and validating the musculoskeletal models that underlie biomechanics-aware control. Kinematics and superficial muscle activity in our model have been validated in Chapter 2 and other studies [58, 181, 243]. However, estimates of metrics such as tendon strains or joint reaction forces remain harder

to verify, given the limited availability of ground-truth measurements. In this sense, medical imaging techniques such as traditional ultrasound or ultrasound elastography appear promising options to validate tendon strains in vivo [244, 245]. On the other hand, sensorized implants [63, 246] have been used to provide benchmarks for joint forces, but it is uncertain how well they capture force behaviour in non-instrumented subjects. However, metrics that remain harder to test experimentally can be made more trustworthy through the validation of other outcomes of the model, since all these quantities emerge from the same underlying physics.

Beyond the shoulder, our approach could be extended to other frequent tendinopathies—such as lateral epicondylitis ("tennis elbow") or patellar tendinitis ("jumper's knee")— where physiotherapy outcomes are still limited by poor understanding of underlying tissue mechanics. Moreover, biomechanics-aware robots could benefit from modeling of the human sensorimotor system, which was not included in this thesis. Including models of reflexive behavior could improve outcomes, providing better understanding of human intentions and reactions to administered therapy.

Robotics as a tool for advancing biomechanics

A particularly exciting prospect is to close the loop in the opposite direction: using robotic systems not only for therapy delivery, but also for improving the biomechanical models themselves. Robots equipped with sensors can act as powerful tools for system identification, capturing precise kinematics, external forces, and even patient-specific responses to exercise. Such data could be used to refine and calibrate models, making them more accurate and individualized. In this sense, rehabilitation robots and exoskeletons may evolve into dual-purpose systems, serving both as therapeutic devices and as experimental platforms for advancing biomechanical knowledge. Improvements in pose estimation—potentially through additional wearable sensors, vision systems, or specialized exoskeleton hardware—could further enhance the fidelity and personalization of models and enable safe assistance in a broader range of movements.

At the same time, it is important to underline that developing excessively personalized models may not scale in practice. Instead, research should identify which level of personalization is needed to deliver meaningful therapy.

In summary, fellow researchers should aim at bridging model fidelity with clinical applicability: validating predictions, engaging patients and PTs in the design of control strategies, and exploiting robotics not only as an assistive technology but also as a scientific instrument. These directions promise to accelerate the transition from laboratory prototypes to clinically adopted systems, bringing biomechanics-aware robotics closer to real impact in healthcare.

7.4. Outlook and broader implications for the future of physiotherapy

At the end of this thesis, I would like to briefly reflect on the broader implications of my work on the future of physiotherapy and, more generally, on physical human-robot interaction. When considering how a rehabilitation robot interacts with a patient, it is evident that many qualities will likely remain uniquely human, most notably the therapist's empathy and emotional connection. Yet there are no fundamental reasons why automated systems could not one day deliver therapy that is at least as effective as, or in certain respects even superior to, that is provided by humans.

The imbalance between the rising demand for rehabilitation and the limited availability of therapists already calls for technologies that can automate parts of the therapeutic process. The potential return on investment is considerable: effective rehabilitation reduces chronic disability and healthcare costs [247–249]. Automating portions of therapy could expand access to timely care, lower the risk of long-term complications, and reduce the burden on insurance systems and public health budgets.

As discussed throughout this thesis, robots equipped with suitable sensors and control algorithms can attain a level of biomechanical awareness beyond what human therapists can directly observe. Their role, however, should not be to replace physiotherapists, but to complement them: filling gaps in biomechanical insight, supporting growing patient numbers, and alleviating the physical fatigue associated with manual therapy. In the long term, we can envision robots as auxiliary tools whose operation is specified at a high level by therapists so that autonomous routines increase therapy time per patient, while therapists focus on supervision and the most complex cases. Instead of replacing PTs, robots should act as their precise apprentices.

Crucially, successful adoption will depend on striking the right balance between technology and therapist education. Human expertise must be preserved and even augmented through these systems, rather than being eroded. Accurate sensing and modeling can play a central role in this balance, unlocking new ways to monitor and guide therapy.

In conclusion, advancements in musculoskeletal modeling are fundamental for assistive robotics, just as robotic applications should not be overlooked by researchers in biomechanics. The potential for cross-pollination and disruptive innovation between the two fields is only beginning to emerge, with implications that extend well beyond rehabilitation toward the design of truly human-centered technologies.

Bibliography

[1] A. T. Nguyen, I. M. Aris, B. D. Snyder, M. B. Harris, J. D. Kang, M. Murray, E. K. Rodriguez and A. Nazarian. 'Musculoskeletal health: an ecological study assessing disease burden and research funding'. In: *The Lancet Regional Health–Americas* 29 (2024).

[2] M. Melchior, Y. Roquelaure, B. Evanoff, J.-F. Chastang, C. Ha, E. Imbernon, M. Goldberg and A. Leclerc. 'Why are manual workers at high risk of upper limb disorders? The role of physical work factors in a random sample of workers in France (the Pays de la Loire study)'. In: *Occupational and environmental medicine* 63.11 (2006), pp. 754–761.

[3] L. Laver, I. P. Pengas and O. Mei-Dan. 'Injuries in extreme sports'. In: *Journal of orthopaedic surgery and research* 12.1 (2017), p. 59.

[4] A. M. Briggs, A. D. Woolf, K. Dreinhöfer, N. Homb, D. G. Hoy, D. Kopansky-Giles, K. Åkesson and L. March. 'Reducing the global burden of musculoskeletal conditions'. In: *Bulletin of the World Health Organization* 96.5 (2018), p. 366.

[5] J. Verhaar, P. Kjærsgaard-Andersen, D. Limb, K.-P. Günther and T. Karachalios. 'The EFORT white book: "orthopaedics and traumatology in Europe"[Internet]'. In: (2021).

[6] A. Cieza, K. Causey, K. Kamenov, S. W. Hanson, S. Chatterji and T. Vos. 'Global estimates of the need for rehabilitation based on the Global Burden of Disease study 2019: a systematic analysis for the Global Burden of Disease Study 2019'. In: *The Lancet* 396.10267 (2020), pp. 2006–2017.

[7] C. Chen, Y. Du, K. Cao, Y. You, L. Pi, D. Jiang, M. Yang, X. Wu, M. Chen, W. Zhou *et al.* 'Global years lived with disability for musculoskeletal disorders in adults 70 Years and older from 1990 to 2019, and projections to 2040'. In: *Heliyon* 10.15 (2024).

[8] M. Morgan, C. Shinost, S. Mendez, J. Klose, G. Lee, R. Forte, T. Lane and M. Klinker. 'The Effect of Early Physical Therapy Intervention on Case Duration and Physical Therapy Visits in Acute Work-Related Musculoskeletal Injuries Across Body Regions: A Retrospective Cohort Study'. In: *JOSPT Open* 3.3 (2025), pp. 1–21.

[9] D. Hertling and R. M. Kessler. *Management of common musculoskeletal disorders: physical therapy principles and methods*. Lippincott Williams & Wilkins, 2006.

[10] M. D. Landry, L. M. Hack, E. Coulson, J. Freburger, M. P. Johnson, R. Katz, J. Kerwin, M. H. Smith, H. C. Wessman, D. G. Venskus *et al.* 'Workforce projections 2010–2020: annual supply and demand forecasting models for physical therapists across the United States'. In: *Physical therapy* 96.1 (2016), pp. 71–80.

[11] P. Zarek, C. Ruttinger, D. Armstrong, R. Chakrabarti, D. R. Hess, T. J. Manal and T. M. Dall. 'Current and Projected Future Supply and Demand for Physical Therapists From 2022 to 2037: A New Approach Using Microsimulation'. In: *Physical Therapy* 105.3 (2025), pzaf014. DOI: 10.1093/ptj/pzaf014. eprint: https://academic.oup.com/ptj/article-pdf/105/3/pzaf014/62213050/pzaf014.pdf. URL: https://doi.org/10.1093/ptj/pzaf014.

[12] J. E. Cromie, V. J. Robertson and M. O. Best. 'Work-related musculoskeletal disorders in physical therapists: prevalence, severity, risks, and responses'. In: *Physical therapy* 80.4 (2000), pp. 336–351.

[13] R. P. Di Fabio. 'Efficacy of manual therapy'. In: *Physical therapy* 72.12 (1992), pp. 853–864.

[14] S. Thomson, C. Jukes and J. Lewis. 'Rehabilitation following surgical repair of the rotator cuff: a systematic review'. In: *Physiotherapy* 102.1 (2016), pp. 20–28.

[15] N. Trofimova. *Photo of an elderly woman*. Accessed: 2025-09-13. 2023. URL: https://unsplash.com/photos/a-woman-walking-down-a-sidewalk-with-a-cane-wiK2icNrdVw?utm_content=creditShareLink&utm_medium=referral&utm_source=unsplash.

[16] okeykat. *Photo of a manual worker*. Accessed: 2025-09-13. 2023. URL: https://unsplash.com/photos/a-person-sweeping-the-street-oVaDLdEQVPw.

[17] E. Chaparro. *Photo of gym exercise*. Accessed: 2025-09-13. 2023. URL: https://unsplash.com/photos/grayscale-photo-of-man-working-out-sHfo3WOgGTU?utm_content=creditShareLink&utm_medium=referral&utm_source=unsplash.

[18] E. Report. *Photo of the Lokomat exoskeleton by Hocoma*. Accessed: 2025-09-13. 2023. URL: https://exoskeletonreport.com/product/lokomat/.

[19] M. Zeman. *Photo of a physiotherapist working*. Accessed: 2025-09-13. 2023. URL: https://unsplash.com/photos/a-woman-getting-a-back-massage-from-a-man-uLWW09LLG04?utm_content=creditShareLink&utm_medium=referral&utm_source=unsplash.

[20] L. Marchal-Crespo and D. J. Reinkensmeyer. 'Review of control strategies for robotic movement training after neurologic injury'. In: *Journal of neuroengineering and rehabilitation* 6 (2009), pp. 1–15.

[21] K. Y. Nam, H. J. Kim, B. S. Kwon, J.-W. Park, H. J. Lee and A. Yoo. 'Robot-assisted gait training (Lokomat) improves walking function and activity in people with spinal cord injury: a systematic review'. In: *Journal of neuroengineering and rehabilitation* 14 (2017), pp. 1–13.

[22] Y. Zimmermann, M. Sommerhalder, P. Wolf, R. Riener and M. Hutter. 'ANYexo 2.0: A fully actuated upper-limb exoskeleton for manipulation and joint-oriented training in all stages of rehabilitation'. In: *IEEE Transactions on Robotics* 39.3 (2023), pp. 2131–2150.

[23] B. Kim and A. D. Deshpande. 'An upper-body rehabilitation exoskeleton Harmony with an anatomical shoulder mechanism: Design, modeling, control, and performance evaluation'. In: *The International Journal of Robotics Research* 36.4 (Apr. 2017), pp. 414–435. ISSN: 0278-3649, 1741-3176. DOI: 10.1177/0278364917706743. (Visited on 21/05/2021).

[24] Hocoma. *ArmeoPower rehabilitation robot*. Accessed: 2025-09-17. URL: https://www.hocoma.com/us/solutions/armeo-power/.

[25] R. Gassert and V. Dietz. 'Rehabilitation robots for the treatment of sensorimotor deficits: a neurophysiological perspective'. In: *Journal of neuroengineering and rehabilitation* 15.1 (2018), p. 46.

[26] E. Basalp, P. Wolf and L. Marchal-Crespo. 'Haptic training: which types facilitate (re) learning of which motor task and for whom? Answers by a review'. In: *IEEE transactions on haptics* 14.4 (2021), pp. 722–739.

[27] A. L. Ratschat, B. M. van Rooij, J. Luijten and L. Marchal-Crespo. 'Evaluating tactile feedback in addition to kinesthetic feedback for haptic shape rendering: a pilot study'. In: *Frontiers in Robotics and AI* 11 (2024), p. 1298537.

[28] S. L. Cucinella, J. C. de Winter, E. Grauwmeijer, M. Evers and L. Marchal-Crespo. 'Towards personalized immersive virtual reality neurorehabilitation: a human-centered design'. In: *Journal of NeuroEngineering and Rehabilitation* 22.1 (2025), p. 7.

[29] A. Van Den Berg, K. L. Poggensee, D. Abbink and L. Marchal-Crespo. 'Visual Disturbances to Avatar Foot Position Increase Step-width Variability in Immersive VR Treadmill Walking'. In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* (2025).

[30] L. Robotics. *ROBERT upper extremity module*. Accessed: 2025-09-17. URL: https://www.lifescience-robotics.com/solutions/upper-extremities/.

[31] 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'. In: *PLoS computational biology* 14.7 (2018), e1006223.

[32] A. Falisse, G. Serrancolí, C. L. Dembia, J. Gillis, I. Jonkers and F. De Groote. ‘Rapid predictive simulations with complex musculoskeletal models suggest that diverse healthy and pathological human gaits can emerge from similar control strategies’. In: *Journal of The Royal Society Interface* 16.157 (2019), p. 20190402.

[33] F. De Groote and A. Falisse. ‘Perspective on musculoskeletal modelling and predictive simulations of human movement to assess the neuromechanics of gait’. In: *Proceedings of the Royal Society B* 288.1946 (2021), p. 20202432.

[34] M. Febrer-Nafria, A. Nasr, M. Ezati, P. Brown, J. M. Font-Llagunes and J. McPhee. ‘Predictive multibody dynamic simulation of human neuromusculoskeletal systems: a review’. In: *Multibody System Dynamics* 58.3 (2023), pp. 299–339.

[35] C. Fang, L. Peternel, A. Seth, M. Sartori, K. Mombaur and E. Yoshida. ‘Human modeling in physical human-robot interaction: A brief survey’. In: *IEEE Robotics and Automation Letters* 8.9 (2023), pp. 5799–5806.

[36] B. Luciani, M. Sommerhalder, M. Gandolla, P. Wolf, F. Braghi and R. Riener. ‘Therapists’ force-profile teach-and-mimic approach for upper-limb rehabilitation exoskeletons’. In: *IEEE Transactions on Medical Robotics and Bionics* (2024).

[37] Y. Tao, Y. Ji, D. Han, H. Gao and T. Wang. ‘A Safe Admittance Boundary Algorithm for Rehabilitation Robot Based on Space Classification Model’. In: *Applied Sciences* 13.9 (2023), p. 5816.

[38] C. Lauretti, F. Cordella, A. L. Ciancio, E. Trigili, J. M. Catalan, F. J. Badesa, S. Crea, S. M. Pagliara, S. Sterzi, N. Vitiello *et al.* ‘Learning by demonstration for motion planning of upper-limb exoskeletons’. In: *Frontiers in neurorobotics* 12 (2018), p. 5.

[39] C. Latella, S. Traversaro, D. Ferigo, Y. Tirupachuri, L. Rapetti, F. J. Andrade Chavez, F. Nori and D. Pucci. ‘Simultaneous floating-base estimation of human kinematics and joint torques’. In: *Sensors* 19.12 (2019), p. 2794.

[40] 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: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2017, pp. 502–507.

[41] 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’. In: *IEEE Robotics and Automation Letters* 7.2 (2022), pp. 4283–4290.

[42] A. Rajagopal, C. L. Dembia, M. S. DeMers, D. D. Delp, J. L. Hicks and S. L. Delp. ‘Full-body musculoskeletal model for muscle-driven simulation of human gait’. In: *IEEE transactions on biomedical engineering* 63.10 (2016), pp. 2068–2079.

[43] E. Beaucage-Gauvreau, W. S. Robertson, S. C. Brandon, R. Fraser, B. J. Freeman, R. B. Graham, D. Thewlis and C. F. Jones. 'Validation of an OpenSim full-body model with detailed lumbar spine for estimating lower lumbar spine loads during symmetric and asymmetric lifting tasks'. In: *Computer methods in biomechanics and biomedical engineering* 22.5 (2019), pp. 451–464.

[44] K. Song, B. M. Gaffney, K. B. Shelburne, C. Pascual-Garrido, J. C. Clohisy and M. D. Harris. 'Dysplastic hip anatomy alters muscle moment arm lengths, lines of action, and contributions to joint reaction forces during gait'. In: *Journal of biomechanics* 110 (2020), p. 109968.

[45] A. Moya-Esteban, G. Durandau, H. Van Der Kooij and M. Sartori. 'Real-time lumbosacral joint loading estimation in exoskeleton-assisted lifting conditions via electromyography-driven musculoskeletal models'. In: *Journal of biomechanics* 157 (2023), p. 111727.

[46] A. Moya-Esteban, S. Sridar, M. I. M. Refai, H. van der Kooij and M. Sartori. 'Adaptive assistance with an active and soft back-support exosuit to unknown external loads via model-based estimates of internal lumbosacral moments'. In: *arXiv preprint arXiv:2311.01843* (2023).

[47] M. Sartori, M. I. Refai, L. A. Gaudio, C. P. Cop, D. Simonetti, F. Damonte, M. Hambly, D. Lloyd, C. Pizzolato and G. Durandau. 'Ceinms-rt: An open-source framework for the continuous neuro-mechanical model-based control of wearable robots'. In: *techrxiv preprint 10.36227/techrxiv.173397962.28177284/v1* (2025).

[48] 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'. In: *Journal of orthopaedics* 10.1 (2013), pp. 8–12.

[49] C. Milgrom, M. Schaffler, S. Gilbert and M. van Holsbeeck. 'Rotator-cuff changes in asymptomatic adults. The effect of age, hand dominance and gender'. In: *The Journal of Bone & Joint Surgery British Volume* 77.2 (1995), pp. 296–298.

[50] R. Ainsworth, J. Lewis and V. Conboy. 'A prospective randomized placebo controlled clinical trial of a rehabilitation programme for patients with a diagnosis of massive rotator cuff tears of the shoulder'. In: *Shoulder & Elbow* 1.1 (2009), pp. 55–60.

[51] B. Mazuquin, M. Moffatt, P. Gill, J. Selfe, J. Rees, S. Drew and C. Littlewood. 'Effectiveness of early versus delayed rehabilitation following rotator cuff repair: systematic review and meta-analyses'. In: *PloS one* 16.5 (2021), e0252137.

[52] C. J. Minns Lowe, J. Moser and K. Barker. 'Living with a symptomatic rotator cuff tear 'bad days, bad nights': a qualitative study'. In: *BMC musculoskeletal disorders* 15.1 (2014), p. 228.

[53] J. K. Mantone, W. Z. Burkhead Jr and J. Noonan Jr. 'Nonoperative treatment of rotator cuff tears'. In: *Orthopedic Clinics of North America* 31.2 (2000), pp. 295–311.

[54] A. Ryösä, K. Laimi, V. Äärimaa, K. Lehtimäki, J. Kukkonen and M. Saltychev. 'Surgery or conservative treatment for rotator cuff tear: a meta-analysis'. In: *Disability and rehabilitation* 39.14 (2017), pp. 1357–1363.

[55] F. V. Sciarretta, D. Moya and K. List. 'Current trends in rehabilitation of rotator cuff injuries'. In: *Sicot-j* 9 (2023).

[56] N. Andarawis-Puri, E. T. Ricchetti and L. J. Soslowsky. 'Rotator cuff tendon strain correlates with tear propagation'. In: *Journal of biomechanics* 42.2 (2009), pp. 158–163.

[57] H. V. Carter. *Anatomical depticion of the human shoulder*. Public domain, via Wikimedia Commons (Accessed: 2025-09-14). 2023. URL: https://commons.wikimedia.org/wiki/File:Gray412.png.

[58] 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'. In: *Frontiers in neurorobotics* 13 (2019), p. 90.

[59] 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'. In: *Plos one* 18.11 (2023), e0295003.

[60] I. Belli, J. M. Prendergast, A. Seth and L. Peternel. 'Biomechanics-Aware Trajectory Optimization for Navigation during Robotic Physiotherapy'. In: *arXiv preprint arXiv:2411.03873* (2024).

[61] 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: *2022 IEEE-RAS 21st International Conference on Humanoid Robots (Humanoids)*. IEEE. 2022, pp. 246–253.

[62] L. Noteboom, I. Belli, M. J. Hoozemans, A. Seth, H. Veeger and F. Van Der Helm. 'Effects of bench press technique variations on musculoskeletal shoulder loads and potential injury risk'. In: *Frontiers in Physiology* 15 (2024), p. 1393235.

[63] I. M. I. Hasan, I. Belli, A. Seth and E. M. Gutierrez-Farewik. 'Modeling glenohumeral stability in musculoskeletal simulations: A validation study with in vivo contact forces'. In: *bioRxiv* (2025), pp. 2025–05.

[64] P. D. Triffitt. 'The relationship between motion of the shoulder and the stated ability to perform activities of daily living'. In: *JBJS* 80.1 (1998), pp. 41–6.

[65] H. Veeger and F. Van Der Helm. 'Shoulder function: the perfect compromise between mobility and stability'. In: *Journal of biomechanics* 40.10 (2007), pp. 2119–2129.

[66] J. E. Labriola, T. Q. Lee, R. E. Debski and P. J. McMahon. 'Stability and instability of the glenohumeral joint: the role of shoulder muscles'. In: *Journal of shoulder and elbow surgery* 14.1 (2005), S32–S38.

[67] S. Lippitt and F. Matsen. 'Mechanisms of glenohumeral joint stability.' In: *Clinical orthopaedics and related research* 291 (1993), pp. 20–28.

[68] T. Yanagawa, C. J. Goodwin, K. B. Shelburne, J. E. Giphart, M. R. Torry and M. G. Pandy. 'Contributions of the individual muscles of the shoulder to glenohumeral joint stability during abduction'. In: *Journal of biomechanical engineering* (2008).

[69] R. Neptune and M. Hull. 'Evaluation of performance criteria for simulation of submaximal steady-state cycling using a forward dynamic model'. In: *Transactions of ASME* (1998).

[70] D. G. Thelen, F. C. Anderson and S. L. Delp. 'Generating dynamic simulations of movement using computed muscle control'. In: *Journal of biomechanics* 36.3 (2003), pp. 321–328.

[71] A. Seth and M. G. Pandy. 'A neuromusculoskeletal tracking method for estimating individual muscle forces in human movement'. In: *Journal of biomechanics* 40.2 (2007), pp. 356–366.

[72] C. Quental, J. Folgado and J. Ambrósio. 'A window moving inverse dynamics optimization for biomechanics of motion'. In: *Multibody System Dynamics* 38.2 (2016), pp. 157–171.

[73] F. De Groote, A. L. Kinney, A. V. Rao and B. J. Fregly. 'Evaluation of direct collocation optimal control problem formulations for solving the muscle redundancy problem'. In: *Annals of biomedical engineering* 44.10 (2016), pp. 2922–2936.

[74] D. E. Hardt. 'Determining muscle forces in the leg during normal human walking—an application and evaluation of optimization methods'. In: *Transactions of the ASME* 100 (1978), pp. 72–78.

[75] R. D. Crowninshield. 'Use of optimization techniques to predict muscle forces'. In: *Transactions of the AMSE* 100 (1978), pp. 88–92.

[76] R. Happee. 'Inverse dynamic optimization including muscular dynamics, a new simulation method applied to goal directed movements'. In: *Journal of biomechanics* 27.7 (1994), pp. 953–960.

[77] M. Ackermann. *Dynamics and energetics of walking with prostheses*. PhD Thesis, 2007.

[78] J. A. Reinbolt. *Static Optimization OpenSim*. 11th Feb. 2023. URL: https://github.com/opensim-org/opensim-core/blob/main/OpenSim/Analyses/StaticOptimization.cpp.

[79] A. Sohane and R. Agarwal. 'Knee muscle force estimating model using machine learning approach'. In: *The Computer Journal* 65.5 (2022), pp. 1167–1177.

[80] A. Nasr, K. A. Inkol, S. Bell and J. McPhee. 'InverseMuscleNET: Alternative Machine Learning Solution to Static Optimization and Inverse Muscle Modeling'. In: *Frontiers in Computational Neuroscience* 15 (2021).

[81] A. Erdemir, S. McLean, W. Herzog and A. J. van den Bogert. 'Model-based estimation of muscle forces exerted during movements'. In: *Clinical biomechanics* 22.2 (2007), pp. 131–154.

[82] M. Ackermann and W. Schiehlen. ‘Physiological methods to solve the force-sharing problem in biomechanics’. In: *Multibody dynamics: Computational methods and applications* (2009), pp. 1–23.

[83] F. E. Zajac. ‘Muscle and tendon: properties, models, scaling, and application to biomechanics and motor control.’ In: *Critical reviews in biomedical engineering* 17.4 (1989), pp. 359–411.

[84] A. Kian, C. Pizzolato, M. Halaki, K. Ginn, D. Lloyd, D. Reed and D. Ackland. ‘Static optimization underestimates antagonist muscle activity at the glenohumeral joint: A musculoskeletal modeling study’. In: *Journal of biomechanics* 97 (2019), p. 109348.

[85] M. H. Akhavanfar, S. C. Brandon, S. H. Brown and R. B. Graham. ‘Development of a novel MATLAB-based framework for implementing mechanical joint stability constraints within OpenSim musculoskeletal models’. In: *Journal of Biomechanics* 91 (2019), pp. 61–68.

[86] S. R. Dubowsky, J. Rasmussen, S. A. Sisto and N. A. Langrana. ‘Validation of a musculoskeletal model of wheelchair propulsion and its application to minimizing shoulder joint forces’. In: *Journal of biomechanics* 41.14 (2008), pp. 2981–2988.

[87] C. Engelhardt, V. Malfroy Camine, D. Ingram, P. Müllhaupt, A. Farron, D. Pioletti and A. Terrier. ‘Comparison of an EMG-based and a stress-based method to predict shoulder muscle forces’. In: *Computer methods in biomechanics and biomedical engineering* 18.12 (2015), pp. 1272–1279.

[88] A. Nikooyan, H. Veeger, P. Westerhoff, B. Bolsterlee, F. Graichen, G. Bergmann and F. Van der Helm. ‘An EMG-driven musculoskeletal model of the shoulder’. In: *Human movement science* 31.2 (2012), pp. 429–447.

[89] F. C. Van der Helm. ‘A finite element musculoskeletal model of the shoulder mechanism’. In: *Journal of biomechanics* 27.5 (1994), pp. 551–569.

[90] F. Steenbrink, J. H. de Groot, H. Veeger, F. van der Helm and P. Rozing. ‘Glenohumeral stability in simulated rotator cuff tears’. In: *Journal of biomechanics* 42.11 (2009), pp. 1740–1745.

[91] I. W. Charlton and G. Johnson. ‘A model for the prediction of the forces at the glenohumeral joint’. In: *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine* 220.8 (2006), pp. 801–812.

[92] D. Blana, J. G. Hincapie, E. K. Chadwick and R. F. Kirsch. ‘A musculoskeletal model of the upper extremity for use in the development of neuroprosthetic systems’. In: *Journal of biomechanics* 41.8 (2008), pp. 1714–1721.

[93] C. R. Dickerson, D. B. Chaffin and R. E. Hughes. ‘A mathematical musculoskeletal shoulder model for proactive ergonomic analysis’. In: *Computer methods in biomechanics and biomedical engineering* 10.6 (2007), pp. 389–400.

[94] C. R. Dickerson, R. E. Hughes and D. B. Chaffin. ‘Experimental evaluation of a computational shoulder musculoskeletal model’. In: *Clinical Biomechanics* 23.7 (2008), pp. 886–894.

[95] D. C. McFarland, A. G. Brynildsen and K. R. Saul. ‘Sensitivity of neuromechanical predictions to choice of glenohumeral stability modeling approach’. In: *Journal of Applied Biomechanics* 36.4 (2020), pp. 249–258.

[96] N. Assila, S. Duprey and M. Begon. ‘Glenohumeral joint and muscles functions during a lifting task’. In: *Journal of biomechanics* 126 (2021), p. 110641.

[97] E. K. Chadwick, D. Blana, R. F. Kirsch and A. J. Van Den Bogert. ‘Real-time simulation of three-dimensional shoulder girdle and arm dynamics’. In: *IEEE Transactions on Biomedical Engineering* 61.7 (2014), pp. 1947–1956.

[98] Y. Blache, M. Begon, B. Michaud, L. Desmoulins, P. Allard and F. Dal Maso. ‘Muscle function in glenohumeral joint stability during lifting task’. In: *PLoS One* 12.12 (2017), e0189406.

[99] 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’. In: *IEEE transactions on biomedical engineering* 54.11 (2007), pp. 1940–1950.

[100] J. L. Hicks, T. K. Uchida, A. Seth, A. Rajagopal and S. L. Delp. ‘Is my model good enough? Best practices for verification and validation of musculoskeletal models and simulations of movement’. In: *Journal of biomechanical engineering* 137.2 (2015).

[101] D. G. Thelen. ‘Adjustment of muscle mechanics model parameters to simulate dynamic contractions in older adults’. In: *J. Biomech. Eng.* 125.1 (2003), pp. 70–77.

[102] A. Seth, R. Matias, A. P. Veloso and S. L. Delp. ‘A biomechanical model of the scapulothoracic joint to accurately capture scapular kinematics during shoulder movements’. In: *PloS one* 11.1 (2016), e0141028.

[103] S. Duprey, F. Billuart, S. Sah, X. Ohl, T. Robert, W. Skalli and X. Wang. ‘Three-dimensional rotations of the scapula during arm abduction: evaluation of the acromion marker cluster method in comparison with a model-based approach using biplanar radiograph images’. In: *Journal of applied biomechanics* 31.5 (2015), pp. 396–402.

[104] T. K. Uchida and A. Seth. ‘Conclusion or Illusion: Quantifying Uncertainty in Inverse Analyses From Marker-Based Motion Capture due to Errors in Marker Registration and Model Scaling’. In: *Frontiers in Bioengineering and Biotechnology* 10 (2022).

[105] T. C. Pataky. ‘One-dimensional statistical parametric mapping in Python’. In: *Computer methods in biomechanics and biomedical engineering* 15.3 (2012), pp. 295–301.

[106] G. M. Sullivan and R. Feinn. ‘Using effect size—or why the P value is not enough’. In: *Journal of graduate medical education* 4.3 (2012), pp. 279–282.

[107] A. Asadi Nikooyan, H. Veeger, E. Chadwick, M. Praagman and F. C. van der Helm. 'Development of a comprehensive musculoskeletal model of the shoulder and elbow'. In: *Medical & biological engineering & computing* 49 (2011), pp. 1425–1435.

[108] W. Wu, P. V. Lee, A. L. Bryant, M. Galea and D. C. Ackland. 'Subject-specific musculoskeletal modeling in the evaluation of shoulder muscle and joint function'. In: *Journal of biomechanics* 49.15 (2016), pp. 3626–3634.

[109] K. R. Holzbaur, W. M. Murray and S. L. Delp. 'A model of the upper extremity for simulating musculoskeletal surgery and analyzing neuromuscular control'. In: *Annals of biomedical engineering* 33 (2005), pp. 829–840.

[110] F. C. Van der Helm, H. Veeger, G. Pronk, L. Van der Woude and R. Rozendal. 'Geometry parameters for musculoskeletal modelling of the shoulder system'. In: *Journal of biomechanics* 25.2 (1992), pp. 129–144.

[111] B. Odle, J. Reinbolt, G. Forrest and T. Dyson-Hudson. 'Construction and evaluation of a model for wheelchair propulsion in an individual with tetraplegia'. In: *Medical & biological engineering & computing* 57.2 (2019), pp. 519–532.

[112] F. C. Anderson and M. G. Pandy. 'Static and dynamic optimization solutions for gait are practically equivalent'. In: *Journal of biomechanics* 34.2 (2001), pp. 153–161.

[113] Q. Carlos, A. Margarida, A. Jorge, G. SB and F. João. 'Influence of the musculotendon dynamics on the muscle force-sharing problem of the shoulder—a fully inverse dynamics approach'. In: *Journal of biomechanical engineering* 140.7 (2018).

[114] M. Żuk, M. Syczewska and C. Pezowicz. 'Influence of uncertainty in selected musculoskeletal model parameters on muscle forces estimated in inverse dynamics-based static optimization and hybrid approach'. In: *Journal of Biomechanical Engineering* 140.12 (2018).

[115] 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'. In: *IEEE Robotics and Automation Letters* 6.4 (2021), pp. 7177–7184. DOI: 10.1109/LRA.2021.3097375.

[116] L. Peternel, C. Fang, N. Tsagarakis and A. Ajoudani. 'A selective muscle fatigue management approach to ergonomic human-robot co-manipulation'. In: *Robotics and Computer-Integrated Manufacturing* 58 (2019), pp. 69–79.

[117] C. Messeri, A. Bicchi, A. M. Zanchettin and P. Rocco. 'A Dynamic Task Allocation Strategy to Mitigate the Human Physical Fatigue in Collaborative Robotics'. In: *IEEE Robotics and Automation Letters* 7.2 (2022), pp. 2178–2185.

[118] A. Leclerc, J. Chastang, I. Niedhammer, M. Landre and Y. Roquelaure. 'Incidence of shoulder pain in repetitive work'. In: *Occupational and environmental medicine* 61.1 (2004), pp. 39–44.

[119] S. T. Seroyer, S. J. Nho, B. R. Bach Jr, C. A. Bush-Joseph, G. P. Nicholson and A. A. Romeo. 'Shoulder pain in the overhead throwing athlete'. In: *Sports health* 1.2 (2009), pp. 108–120.

[120] J. L. Codding and J. D. Keener. 'Natural history of degenerative rotator cuff tears'. In: *Current reviews in musculoskeletal medicine* 11 (2018), pp. 77–85.

[121] J. L. Zimbelman, S. P. Juraschek, X. Zhang and V. W.-H. Lin. 'Physical therapy workforce in the United States: forecasting nationwide shortages'. In: *PM&R* 2.11 (2010), pp. 1021–1029.

[122] B. Østerås, H. Østerås, T. A. Torstensen and O. Vasseljen. 'Dose–response effects of medical exercise therapy in patients with patellofemoral pain syndrome: a randomised controlled clinical trial'. In: *Physiotherapy* 99.2 (2013), pp. 126–131.

[123] T. Proietti, V. Crocher, A. Roby-Brami and N. Jarrasse. 'Upper-limb robotic exoskeletons for neurorehabilitation: a review on control strategies'. In: *IEEE reviews in biomedical engineering* 9 (2016), pp. 4–14.

[124] M. Sommerhalder, Y. Zimmermann, J. Song, R. Riener and P. Wolf. 'Polymorphic Control Framework for Automated and Individualized Robot-Assisted Rehabilitation'. In: *IEEE Transactions on Robotics* (2023).

[125] D. Calafiore, F. Negrini, N. Tottoli, F. Ferraro, O. Ozyemisci-Taskiran *et al.* 'Efficacy of robotic exoskeleton for gait rehabilitation in patients with subacute stroke: a systematic review'. In: *European Journal of Physical and Rehabilitation Medicine* 58.1 (2022), p. 1.

[126] P. K. Jamwal, S. Hussain and M. H. Ghayesh. 'Robotic orthoses for gait rehabilitation: An overview of mechanical design and control strategies'. In: *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine* 234.5 (2020), pp. 444–457.

[127] K. L. Poggensee and S. H. Collins. 'How adaptation, training, and customization contribute to benefits from exoskeleton assistance'. In: *Science Robotics* 6.58 (2021), eabf1078.

[128] Y. Cen, J. Yuan, S. Ma, J. Luo and H. Wang. 'Trajectory Optimization Algorithm of Trajectory Rehabilitation Training Mode for Rehabilitation Robot'. In: *2022 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. IEEE. 2022, pp. 2153–2158.

[129] B. Luciani, L. Roveda, F. Braghin, A. Pedrocchi and M. Gandolla. 'Trajectory learning by therapists' demonstrations for an upper limb rehabilitation exoskeleton'. In: *IEEE Robotics and Automation Letters* (2023).

[130] C. Fang, L. Peternel, A. Seth, M. Sartori, K. Mombaur and E. Yoshida. 'Human Modeling in Physical Human-Robot Interaction: A Brief Survey'. In: *IEEE Robotics and Automation Letters* (2023).

[131] M. G. Carmichael and D. Liu. 'Estimating physical assistance need using a musculoskeletal model'. In: *IEEE Transactions on Biomedical Engineering* 60.7 (2013), pp. 1912–1919.

[132] 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: *IECON 2019-45th Annual Conference of the IEEE Industrial Electronics Society*. Vol. 1. IEEE. 2019, pp. 5394–5400.

[133] J. Fang and Y. Yuan. ‘Human-in-the-loop optimization of wearable robots to reduce the human metabolic energy cost in physical movements’. In: *Robotics and Autonomous Systems* 127 (2020), p. 103495.

[134] D. F. Gordon, C. McGreavy, A. Christou and S. Vijayakumar. ‘Human-in-the-loop optimization of exoskeleton assistance via online simulation of metabolic cost’. In: *IEEE Transactions on Robotics* 38.3 (2022), pp. 1410–1429.

[135] M. Gallois, M. Manzano, S. Guégan, N. Vignais, M. Babel and C. Pontonnier. ‘Effort generation capabilites mapping for personalized robotic assistance of the elbow’. In: *IEEE International Conference on Rehabilitation Robotics (ICORR)*. 2025.

[136] D. G. Thelen, F. C. Anderson and S. L. Delp. ‘Generating dynamic simulations of movement using computed muscle control’. In: *Journal of biomechanics* 36.3 (2003), pp. 321–328.

[137] C. L. Dembia, N. A. Bianco, A. Falisse, J. L. Hicks and S. L. Delp. ‘Opensim moco: Musculoskeletal optimal control’. In: *PLOS Computational Biology* 16.12 (2020), e1008493.

[138] E. Peiros, Z.-Y. Chiu, Y. Zhi, N. Shinde and M. C. Yip. ‘Finding biomechanically safe trajectories for robot manipulation of the human body in a search and rescue scenario’. In: *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2023, pp. 167–173.

[139] V. Vatsal and G. Hoffman. ‘Biomechanical motion planning for a wearable robotic forearm’. In: *IEEE Robotics and Automation Letters* 6.3 (2021), pp. 5024–5031.

[140] P. K. Jamwal, S. Hussain, Y. H. Tsoi and S. Q. Xie. ‘Musculoskeletal model for path generation and modification of an ankle rehabilitation robot’. In: *IEEE Transactions on Human-Machine Systems* 50.5 (2020), pp. 373–383.

[141] C. Latella, Y. Tirupachuri, L. Tagliapietra, L. Rapetti, B. Schirrmeister, J. Bornmann, D. Gorjan, J. Čamernik, P. Maurice, L. Fritzsche *et al.* ‘Analysis of human whole-body joint torques during overhead work with a passive exoskeleton’. In: *IEEE Transactions on Human-Machine Systems* 52.5 (2021), pp. 1060–1068.

[142] X. Bao, Z. Sheng, B. E. Dicianno and N. Sharma. ‘A tube-based model predictive control method to regulate a knee joint with functional electrical stimulation and electric motor assist’. In: *IEEE Transactions on Control Systems Technology* 29.5 (2020), pp. 2180–2191.

[143] G. Clark and H. B. Amor. ‘Learning ergonomic control in human–robot symbiotic walking’. In: *IEEE Transactions on Robotics* 39.1 (2022), pp. 327–342.

[144] 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'. In: *IEEE Robotics and Automation Letters* 6.4 (2021), pp. 7177–7184.

[145] Y. Chen, F. Jiang, H. Li, S. Chen, Y. Qiao, Y. Li, Y. Hua, J. Chen and Y. Ge. 'Retears and concomitant functional impairments after rotator cuff repair: shoulder activity as a risk factor'. In: *The American Journal of Sports Medicine* 48.4 (2020), pp. 931–938.

[146] R. Drillis, R. Contini and M. Bluestein. 'Body segment parameters'. In: *Artificial limbs* 8.1 (1964), pp. 44–66.

[147] 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: *2023 IEEE-RAS 22nd International Conference on Humanoid Robots (Humanoids)*. IEEE. 2023, pp. 1–8.

[148] 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'. In: *PLoS One* 14.10 (2019), e0217730.

[149] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings and M. Diehl. 'CasADi – A software framework for nonlinear optimization and optimal control'. In: *Mathematical Programming Computation* 11.1 (2019), pp. 1–36. DOI: 10.1007/s12532-018-0139-4.

[150] 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'. In: *hal-01615132* (2017).

[151] A. Wächter and L. T. Biegler. 'On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming'. In: *Mathematical programming* 106 (2006), pp. 25–57.

[152] I. S. Duff and J. K. Reid. 'The multifrontal solution of indefinite sparse symmetric linear'. In: *ACM Transactions on Mathematical Software (TOMS)* 9.3 (1983), pp. 302–325.

[153] G. C. Terry and T. M. Chopp. 'Functional anatomy of the shoulder'. In: *Journal of athletic training* 35.3 (2000), p. 248.

[154] T. Flash and N. Hogan. 'The coordination of arm movements: an experimentally confirmed mathematical model'. In: *Journal of neuroscience* 5.7 (1985), pp. 1688–1703.

[155] S. P. Sitole and F. C. Sup. 'Continuous Prediction of Human Joint Mechanics Using EMG Signals: A Review of Model-Based and Model-Free Approaches'. In: *IEEE Transactions on Medical Robotics and Bionics* 5.3 (2023), pp. 528–546. DOI: 10.1109/TMRB.2023.3292451.

[156] W. Mugge, D. A. Abbink, A. C. Schouten, J. P. Dewald and F. C. Van Der Helm. 'A rigorous model of reflex function indicates that position and force feedback are flexibly tuned to position and force tasks'. In: *Experimental brain research* 200 (2010), pp. 325–340.

[157] P. C. Lastayo, T. Wright, R. Jaffe and J. Hartzel. 'Continuous passive motion after repair of the rotator cuff. A prospective outcome study'. In: *JBJS* 80.7 (1998), pp. 1002–11.

[158] M. Manzano, S. Guégan, R. Le Breton, L. Devigne and M. Babel. 'Model-based upper-limb gravity compensation strategies for active dynamic arm supports'. In: *2023 International Conference on Rehabilitation Robotics (ICORR)*. IEEE. 2023, pp. 1–6.

[159] M. Urwin, D. Symmons, T. Allison, T. Brammah, H. Busby, M. Roxby, A. Simmons and G. Williams. 'Estimating the burden of musculoskeletal disorders in the community: the comparative prevalence of symptoms at different anatomical sites, and the relation to social deprivation'. In: *Annals of the Rheumatic Diseases* 57.11 (1998), pp. 649–655. ISSN: 0003-4967. DOI: 10.1136/ard.57.11.649.

[160] M. Khan and J. J. P. Warner. 'Cochrane in CORR ®: Manual therapy and exercise for rotator cuff disease'. In: *Clin. Orthop. Relat. Res.* 475.7 (July 2017), pp. 1779–1785.

[161] B. Mazuquin, M. Moffatt, P. Gill, J. Selfe, J. Rees, S. Drew and C. Littlewood. 'Effectiveness of early versus delayed rehabilitation following rotator cuff repair: Systematic review and meta-analyses'. In: *PLoS One* 16.5 (May 2021), e0252137.

[162] 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'. In: *J. Orthop.* 10.1 (Feb. 2013), pp. 8–12.

[163] C. Milgrom, M. Schaffler, S. Gilbert and M. van Holsbeeck. 'Rotator-cuff changes in asymptomatic adults. The effect of age, hand dominance and gender'. In: *J. Bone Joint Surg. Br.* 77.2 (Mar. 1995), pp. 296–298.

[164] F. Jobe, R. Kvitne and C. Giangarra. 'Shoulder pain in the overhand or throwing athlete. The relationship of anterior instability and rotator cuff impingement'. In: *Orthopaedic review* 18.9 (Sept. 1989), pp. 963–975. ISSN: 0094-6591.

[165] G. K. Singh, S. Srivastava, M. Kumar and S. Ratnakar. 'Effects of selected rehabilitative exercises on external rotator muscles and trapezius muscles of masonry workers'. In: *Work* 60.3 (2018), pp. 437–444.

[166] S. T. Seroyer, S. J. Nho, B. R. Bach Jr, C. A. Bush-Joseph, G. P. Nicholson and A. A. Romeo. 'Shoulder pain in the overhead throwing athlete'. In: *Sports Health* 1.2 (Mar. 2009), pp. 108–120.

[167] U. G. Longo, A. Berton, L. Risi Ambrogioni, D. Lo Presti, A. Carnevale, V. Candela, G. Stelitano, E. Schena, A. Nazarian and V. Denaro. 'Cost-effectiveness of supervised versus unsupervised rehabilitation for rotator-cuff repair: Systematic review and meta-analysis'. In: *Int. J. Environ. Res. Public Health* 17.8 (Apr. 2020), p. 2852.

[168] C. Littlewood, P. Malliaras and K. Chance-Larsen. 'Therapeutic exercise for rotator cuff tendinopathy: a systematic review of contextual factors and prescription parameters'. In: *Int. J. Rehabil. Res.* 38.2 (June 2015), pp. 95–106.

[169] J. Kukkonen, A. Joukainen, J. Lehtinen, K. T. Mattila, E. K. J. Tuominen, T. Kauko and V. Aärimaa. 'Treatment of non-traumatic rotator cuff tears: A randomised controlled trial with one-year clinical results'. In: *Bone Joint J.* 96-B.1 (Jan. 2014), pp. 75–81.

[170] P. C. Lastayo, T. Wright, R. Jaffe and J. Hartzel. 'Continuous passive motion after repair of the rotator cuff. A prospective outcome study'. In: *J. Bone Joint Surg. Am.* 80.7 (July 1998), pp. 1002–1011.

[171] B. Østerås, H. Østerås, T. A. Torstensen and O. Vasseljen. 'Dose-response effects of medical exercise therapy in patients with patellofemoral pain syndrome: a randomised controlled clinical trial'. In: *Physiotherapy* 99.2 (June 2013), pp. 126–131.

[172] U. G. Longo, L. R. Ambrogioni, A. Berton, V. Candela, F. Migliorini, A. Carnevale, E. Schena, A. Nazarian, J. DeAngelis and V. Denaro. 'Conservative versus accelerated rehabilitation after rotator cuff repair: a systematic review and meta-analysis'. In: *BMC Musculoskeletal Disorders* 22.1 (July 2021). DOI: 10.1186/s12891-021-04397-0. URL: https://doi.org/10.1186/s12891-021-04397-0.

[173] G. Prange, M. Jannink, C. Groothuis-Oudshoorn, H. Hermens and M. IJzerman. 'Systematic review of the effect of robot-aided therapy on recovery of the hemiparetic arm after stroke'. In: *Journal of rehabilitation research and development* 43 (Mar. 2006), pp. 171–84. DOI: 10.1682/JRRD.2005.04.0076.

[174] K. Yamaguchi, A. Tetro, O. Blam, B. A. Evanoff, S. A. Teefey and W. D. Middleton. 'Natural history of asymptomatic rotator cuff tears: A longitudinal analysis of asymptomatic tears detected sonographically'. In: *Journal of Shoulder and Elbow Surgery* 10.3 (2001), pp. 199–203. ISSN: 1058-2746. DOI: https://doi.org/10.1067/mse.2001.113086.

[175] A. S. Niyetkaliyev, S. Hussain, M. H. Ghayesh and G. Alici. 'Review on design and control aspects of robotic shoulder rehabilitation orthoses'. In: *IEEE Transactions on Human-Machine Systems* 47.6 (2017), pp. 1134–1145.

[176] S. Dalla Gasperina, L. Roveda, A. Pedrocchi, F. Braghin and M. Gandolla. 'Review on patient-cooperative control strategies for upper-limb rehabilitation exoskeletons'. In: *Frontiers in Robotics and AI* 8 (2021), p. 745018.

[177] R. Fareh, A. Elsabe, M. Baziyad, T. Kawser, B. Brahmi and M. H. Rahman. 'Will Your Next Therapist Be a Robot?—A Review of the Advancements in Robotic Upper Extremity Rehabilitation'. In: *Sensors* 23.11 (2023), p. 5054.

[178] S. Buccelli, F. Tessari, F. Fanin, L. De Guglielmo, G. Capitta, C. Piezzo, A. Bruschi, F. Van Son, S. Scarpetta, A. Succi *et al.* 'A gravity-compensated upper-limb exoskeleton for functional rehabilitation of the shoulder complex'. In: *Applied Sciences* 12.7 (2022), p. 3364.

[179] M. Tröbinger, A. Costinescu, H. Xing, J. Elsner, T. Hu, A. Naceri, L. Figueredo, E. Jensen, D. Burschka and S. Haddadin. 'A Dual Doctor-Patient Twin Paradigm for Transparent Remote Examination, Diagnosis, and Rehabilitation'. In: *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2021, pp. 2933–2940.

[180] S. Masiero, A. Celia, G. Rosati and M. Armani. 'Robotic-Assisted Rehabilitation of the Upper Limb After Acute Stroke'. In: *Archives of Physical Medicine and Rehabilitation* 88.2 (2007), pp. 142–149. ISSN: 0003-9993. DOI: https://doi.org/10.1016/j.apmr.2006.10.032.

[181] A. Seth, R. Matias, A. P. Veloso and S. L. Delp. 'A biomechanical model of the scapulothoracic joint to accurately capture scapular kinematics during shoulder movements'. In: *PloS one* 11.1 (2016), e0141028.

[182] 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'. In: *IEEE transactions on biomedical engineering* 54.11 (2007), pp. 1940–1950.

[183] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, İ. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt and SciPy 1.0 Contributors. 'SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python'. In: *Nature Methods* 17 (2020), pp. 261–272. DOI: 10.1038/s41592-019-0686-2.

[184] A. De Luca, A. Albu-Schaffer, S. Haddadin and G. Hirzinger. 'Collision detection and safe reaction with the DLR-III lightweight manipulator arm'. In: *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE. 2006, pp. 1623–1630.

[185] A. Albu-Schaffer, C. Ott, U. Frese and G. Hirzinger. 'Cartesian impedance control of redundant robots: Recent results with the DLR-light-weight-arms'. In: *2003 IEEE International Conference on Robotics and Automation*. Vol. 3. 2003, pp. 3704–3709.

[186] D. A. Winter. *Biomechanics and motor control of human movement*. 4th ed. Chichester, England: John Wiley & Sons, Sept. 2009.

[187] I. Belli, Y. Hu, J. M. Prendergast, D. Abbink, A. Seth and L. Peternel. 'A Shared Control Approach to Safely Limiting Patient Motion Based on Tendon Strain During Robotic-Assisted Shoulder Rehabilitation'. In: *2025 International Conference On Rehabilitation Robotics (ICORR)*. IEEE. 2025, pp. 1071–1077.

[188] T. A. Sgroi and M. Cilenti. 'Rotator cuff repair: post-operative rehabilitation concepts'. In: *Current reviews in musculoskeletal medicine* 11 (2018), pp. 86–91.

[189] 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'. In: *IEEE Transactions on Human-Machine Systems* 48.5 (2018), pp. 509–525.

[190] 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'. In: *Applied Mechanics Reviews* 70.1 (2018), p. 010804.

[191] N. Hogan. 'Impedance control - An approach to manipulation. I Theory. II - Implementation. III - Applications'. In: *ASME Transactions Journal of Dynamic Systems and Measurement Control B* 107 (1985), pp. 1–24.

[192] J. Ravenberg, I. Belli, J. M. Prendergast, A. Seth and L. Peternel. 'Creating Discomfort Maps via Hand-held Human Feedback Interface for Robotic Shoulder Physiotherapy'. In: *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2024, pp. 4664–4671.

[193] 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'. In: *Journal of Orthopaedics* 10.1 (Mar. 2013), pp. 8–12. ISSN: 0972-978X. DOI: 10.1016/j.jor.2013.01.008. URL: https://www.sciencedirect.com/science/article/pii/S0972978X13000093 (visited on 20/12/2023).

[194] T. Proietti, V. Crocher, A. Roby-Brami and N. Jarrasse. 'Upper-Limb Robotic Exoskeletons for Neurorehabilitation: A Review on Control Strategies'. eng. In: *IEEE reviews in biomedical engineering* 9 (2016), pp. 4–14. ISSN: 1941-1189. DOI: 10.1109/RBME.2016.2552201.

[195] H. Østerås and T. A. Torstensen. 'The Dose-Response Effect of Medical Exercise Therapy on Impairment in Patients with Unilateral Longstanding Subacromial Pain'. In: *The Open Orthopaedics Journal* 4 (Jan. 2010), pp. 1–6. ISSN: 1874-3250. DOI: 10.2174/1874325001004010001. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2817873/ (visited on 20/12/2023).

[196] 'Ticking timebomb: Without immediate action, health and care workforce gaps in the European Region could spell disaster'. In: *World Health Organization* (Sept. 2022). URL: https://www.who.int/europe/news/item/14-09-2022-ticking-timebomb--without-immediate-action--health-and-care-workforce-gaps-in-the-european-region-could-spell-disaster.

[197] D. Simonetti, L. Zollo, E. Papaleo, G. Carpino and E. Guglielmelli. 'Multimodal adaptive interfaces for 3D robot-mediated upper limb neuro-rehabilitation: An overview of bio-cooperative systems'. In: *Robotics and Autonomous Systems* 85 (Nov. 2016), pp. 62–72. ISSN: 0921-8890. DOI: 10.1016/j.robot.2016.08.012. URL: https://www.sciencedirect.com/science/article/pii/S0921889016304808 (visited on 20/12/2023).

[198] F. Scotto di Luzio, D. Simonetti, F. Cordella, S. Miccinilli, S. Sterzi, F. Draicchio and L. Zollo. 'Bio-Cooperative Approach for the Human-in-the-Loop Control of an End-Effector Rehabilitation Robot'. In: *Frontiers in Neurorobotics* 12 (2018). ISSN: 1662-5218. URL: https://www.frontiersin.org/articles/10.3389/fnbot.2018.00067 (visited on 20/12/2023).

[199] W. Kim, J. Lee, L. Peternel, N. Tsagarakis and A. Ajoudani. 'Anticipatory Robot Assistance for the Prevention of Human Static Joint Overloading in Human-Robot Collaboration'. In: *IEEE Robotics and Automation Letters* PP (July 2017), pp. 1–1. DOI: 10.1109/LRA.2017.2729666.

[200] L. Peternel, C. Fang, N. Tsagarakis and A. Ajoudani. 'A selective muscle fatigue management approach to ergonomic human-robot co-manipulation'. In: *Robotics and Computer-Integrated Manufacturing* 58 (Aug. 2019), pp. 69–79. DOI: 10.1016/j.rcim.2019.01.013.

[201] L. F. C. Figueredo, R. C. Aguiar, L. Chen, S. Chakrabarty, M. R. Dogar and A. G. Cohn. 'Human Comfortability: Integrating Ergonomics and Muscular-Informed Metrics for Manipulability Analysis During Human-Robot Collaboration'. en. In: *IEEE Robotics and Automation Letters* 6.2 (Apr. 2021), pp. 351–358. ISSN: 2377-3766, 2377-3774. DOI: 10.1109/LRA.2020.3043173. URL: https://ieeexplore.ieee.org/document/9286730/ (visited on 20/12/2023).

[202] T. Petrič, L. Peternel, J. Morimoto and J. Babič. 'Assistive arm-exoskeleton control based on human muscular manipulability'. English. In: *Frontiers in Neurorobotics* 13 (2019). ISSN: 1662-5218. DOI: 10.3389/fnbot.2019.00030.

[203] J. Prendergast, S. Balvert, T. Driessen, A. Seth and L. Peternel. 'Biomechanics Aware Collaborative Robot System for Delivery of Safe Physical Therapy in Shoulder Rehabilitation'. English. In: *IEEE Robotics and Automation Letters* 6.4 (2021), pp. 7177–7184. ISSN: 2377-3766. DOI: 10.1109/LRA.2021.3097375.

[204] S. Balvert, M. Prendergast, I. Belli, A. Seth and L. Peternel. 'Enabling Patient- and Teleoperator-led Robotic Physiotherapy via Strain Map Segmentation and Shared-authority'. In: Nov. 2022. DOI: 10.1109/Humanoids53995.2022.10000234.

[205] R. B. Fillingim. 'Individual differences in pain: understanding the mosaic that makes pain personal'. en-US. In: *PAIN* 158 (Apr. 2017), S11. ISSN: 0304-3959. DOI: 10.1097/j.pain.0000000000000775. URL: https://journals.lww.com/pain/fulltext/2017/04001/individual_

differences_in_pain__understanding_the.3.aspx (visited on 20/12/2023).

[206] R. Cowen, M. K. Stasiowska, H. Laycock and C. Bantel. 'Assessing pain objectively: the use of physiological markers'. In: *Anaesthesia* 70.7 (July 2015). Publisher: John Wiley & Sons, Ltd, pp. 828–847. ISSN: 0003-2409. DOI: 10.1111/anae.13018. URL: https://associationofanaesthetists.publications.onlinelibrary.wiley.com/doi/10.1111/anae.13018 (visited on 20/12/2023).

[207] G. D. De Sario, C. R. Haider, K. C. Maita, R. A. Torres-Guzman, O. S. Emam, F. R. Avila, J. P. Garcia, S. Borna, C. J. McLeod, C. J. Bruce, R. E. Carter and A. J. Forte. 'Using AI to Detect Pain through Facial Expressions: A Review'. In: *Bioengineering* 10.5 (May 2023), p. 548. ISSN: 2306-5354. DOI: 10.3390/bioengineering10050548. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10215219/ (visited on 20/12/2023).

[208] E. M. Boormans, P. J. Van Kesteren, R. S. Perez, H. A. Brölmann and W. W. Zuurmond. 'Reliability of a Continuous Pain Score Meter: Real Time Pain Measurement'. en. In: *Pain Practice* 9.2 (2009). _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1533-2500.2009.00260.x, pp. 100–104. ISSN: 1533-2500. DOI: 10.1111/j.1533-2500.2009.00260.x. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1533-2500.2009.00260.x (visited on 20/12/2023).

[209] A. van Wijk, F. Lobbezoo and J. Hoogstraten. 'Reliability and validity of a continuous pain registration procedure'. en. In: *European Journal of Pain* 17.3 (2013). _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/j.1532-2149.2012.00194.x, pp. 394–401. ISSN: 1532-2149. DOI: 10.1002/j.1532-2149.2012.00194.x. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/j.1532-2149.2012.00194.x (visited on 20/12/2023).

[210] N. Schaffner, G. Folkers, S. Käppeli, M. Musholt, G. F. L. Hofbauer and V. Candia. 'A New Tool for Real-Time Pain Assessment in Experimental and Clinical Environments'. In: *PLoS ONE* 7.11 (Nov. 2012), e51014. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0051014. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3511427/ (visited on 20/12/2023).

[211] D. M. Böing-Meßing, F. Tomschi, T. Cegla and T. Hilberg. 'The eEgg: Evaluation of a New Device to Measure Pain'. In: *Frontiers in Physiology* 13 (2022). ISSN: 1664-042X. URL: https://www.frontiersin.org/articles/10.3389/fphys.2022.832172 (visited on 03/04/2023).

[212] I. S. K. Thong, M. P. Jensen, J. Miró and G. Tan. 'The validity of pain intensity measures: what do the NRS, VAS, VRS, and FPS-R measure?' en. In: *Scandinavian Journal of Pain* 18.1 (Jan. 2018). Publisher: De Gruyter, pp. 99–107. ISSN: 1877-8879. DOI: 10.1515/sjpain-2018-0012. URL:

https://www.degruyter.com/document/doi/10.1515/sjpain-2018-0012/html (visited on 29/03/2023).

[213] T. Paolucci, F. Agostini, M. Mangone, A. Bernetti, L. Pezzi, V. Liotti, E. Recubini, C. Cantarella, R. G. Bellomo, C. D'Aurizio *et al.* 'Robotic rehabilitation for end-effector device and botulinum toxin in upper limb rehabilitation in chronic post-stroke patients: an integrated rehabilitative approach'. In: *Neurological Sciences* (2021), pp. 1–11.

[214] Y. Bouteraa, I. B. Abdallah, K. Alnowaiser and A. Ibrahim. 'Smart solution for pain detection in remote rehabilitation'. In: *Alexandria Engineering Journal* 60.4 (2021), pp. 3485–3500.

[215] L. Dai, J. Broekens and K. P. Truong. 'Real-time pain detection in facial expressions for health robotics'. In: *2019 8th International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW)*. IEEE. 2019, pp. 277–283.

[216] L. Peternel, T. Petrič and J. Babič. 'Robotic assembly solution by human-in-the-loop teaching method based on real-time stiffness modulation'. English. In: *Autonomous Robots* 42.1 (2018), pp. 1–17. ISSN: 0929-5593. DOI: 10.1007/s10514-017-9635-z.

[217] A. Albu-Schaffer, C. Ott, U. Frese and G. Hirzinger. 'Cartesian impedance control of redundant robots: recent results with the DLR-light-weight-arms'. In: *2003 IEEE International Conference on Robotics and Automation (Cat. No.03CH37422)*. Vol. 3. ISSN: 1050-4729. Sept. 2003, 3704–3709 vol.3. DOI: 10.1109/ROBOT.2003.1242165.

[218] C. C. De Wit, B. Siciliano and G. Bastin, eds. *Theory of Robot Control*. en. Communications and Control Engineering. London: Springer, 1996. ISBN: 978-1-4471-1503-8 978-1-4471-1501-4. DOI: 10.1007/978-1-4471-1501-4. URL: http://link.springer.com/10.1007/978-1-4471-1501-4 (visited on 01/01/2024).

[219] I. Jammeli, A. Chemori, H. Moon, S. Elloumi and S. Mohammed. 'An assistive explicit model predictive control framework for a knee rehabilitation exoskeleton'. In: *IEEE/ASME Transactions on Mechatronics* 27.5 (2021), pp. 3636–3647.

[220] G. Bertoni, G. Leuzzi, M. Job, M. De Simone and M. Testa. 'Exploring knowledge, perception, and use of surface electromyography in physiotherapy post graduate trainees in Italy: a single center preliminary survey'. In: *Frontiers in rehabilitation sciences* 5 (2024), p. 1489927.

[221] L. McManus, G. De Vito and M. M. Lowery. 'Analysis and biophysics of surface EMG for physiotherapists and kinesiologists: Toward a common language with rehabilitation engineers'. In: *Frontiers in neurology* 11 (2020), p. 576729.

[222] L. Peternel, N. Tsagarakis, D. Caldwell and A. Ajoudani. 'Robot adaptation to human physical fatigue in human–robot co-manipulation'. In: *Autonomous Robots* 42.5 (2018), pp. 1011–1021.

[223] L. Peternel, T. Noda, T. Petrič, A. Ude, J. Morimoto and J. Babič. 'Adaptive control of exoskeleton robots for periodic assistive behaviours based on EMG feedback minimisation'. In: *PloS one* 11.2 (2016), e0148942.

[224] M. Gazzoni, B. Afsharipour and R. Merletti. 'Surface EMG in ergonomics and occupational medicine'. In: *Surface electromyography: physiology, engineering, and applications* (2016), pp. 361–391.

[225] S. Kumar. 'Electromyography in ergonomics'. In: *Electromyography in ergonomics*. Routledge, 2017, pp. 1–50.

[226] T. McFarland and S. Fischer. 'Considerations for industrial use: a systematic review of the impact of active and passive upper limb exoskeletons on physical exposures'. In: *IISE Transactions on Occupational Ergonomics and Human Factors* 7.3-4 (2019), pp. 322–347.

[227] H. Fritz, D. Patzer and S. S. Galen. 'Robotic exoskeletons for reengaging in everyday activities: promises, pitfalls, and opportunities'. In: *Disability and rehabilitation* 41.5 (2019), pp. 560–563.

[228] T. Van Wouwe, L. H. Ting and F. De Groote. 'An approximate stochastic optimal control framework to simulate nonlinear neuro-musculoskeletal models in the presence of noise'. In: *PLoS computational biology* 18.6 (2022), e1009338.

[229] P. Lele, D. Sinclair and G. Weddell. 'The reaction time to touch'. In: *The Journal of physiology* 123.1 (1954), p. 187.

[230] V. Caggiano, H. Wang, G. Durandau, M. Sartori and V. Kumar. 'MyoSuite–A contact-rich simulation suite for musculoskeletal motor control'. In: *arXiv preprint arXiv:2205.13600* (2022).

[231] G. Williams, A. Aldrich and E. A. Theodorou. 'Model predictive path integral control: From theory to parallel computation'. In: *Journal of Guidance, Control, and Dynamics* 40.2 (2017), pp. 344–357.

[232] M. Kazim, J. Hong, M.-G. Kim and K.-K. K. Kim. 'Recent advances in path integral control for trajectory optimization: An overview in theoretical and algorithmic perspectives'. In: *Annual Reviews in Control* 57 (2024), p. 100931.

[233] E. AG. *FORCESPRO*. https://forces.embotech.com/. 2014.

[234] N. Hogan. 'Impedance control - An approach to manipulation. I - Theory. II - Implementation. III - Applications'. In: *ASME Transactions Journal of Dynamic Systems and Measurement Control B* 107 (Mar. 1985), pp. 1–24.

[235] L. Song, C. Ju, H. Cui, Y. Qu, X. Xu and C. Chen. 'Research on Control Strategy Technology of Upper Limb Exoskeleton Robots'. In: *Machines* 13.3 (2025), p. 207.

[236] A. Q. Keemink, H. Van der Kooij and A. H. Stienen. 'Admittance control for physical human–robot interaction'. In: *The International Journal of Robotics Research* 37.11 (2018), pp. 1421–1444.

[237] R. J. Escarabajal, E. París, M. Jamšek, T. Petrič, Á. Valera, V. Mata and J. Babič. 'Muscle-Targeted Robotic Assistance and Augmentation of Human Motion'. In: *IEEE Transactions on Cybernetics* (2025).

[238] F. Zahedi and H. Lee. ‘Biomechanics-Based User-Adaptive Variable Impedance Control for Enhanced Physical Human–Robot Interaction Using Bayesian Optimization’. In: *Advanced Intelligent Systems* 7.2 (2025), p. 2400333.

[239] K. P. Wabersich and M. N. Zeilinger. ‘A predictive safety filter for learning-based control of constrained nonlinear dynamical systems’. In: *Automatica* 129 (2021), p. 109597.

[240] T. Van Wouwe, S. Lee, A. Falisse, S. Delp and C. K. Liu. ‘Diffusionposer: Real-time human motion reconstruction from arbitrary sparse sensors using autoregressive diffusion’. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2024, pp. 2513–2523.

[241] C. Celemin, R. Pérez-Dattari, E. Chisari, G. Franzese, L. de Souza Rosa, R. Prakash, Z. Ajanović, M. Ferraz, A. Valada, J. Kober *et al.* ‘Interactive imitation learning in robotics: A survey’. In: *Foundations and Trends® in Robotics* 10.1-2 (2022), pp. 1–197.

[242] E. Halilaj, A. Rajagopal, M. Fiterau, J. L. Hicks, T. J. Hastie and S. L. Delp. ‘Machine learning in human movement biomechanics: Best practices, common pitfalls, and new opportunities’. In: *Journal of biomechanics* 81 (2018), pp. 1–11.

[243] A. G. C. S. Carvalho, A. Seth, I. Belli, W. de Vries and H. Veeger. ‘Muscle contributions to work during manual wheelchair propulsion’. In: (2024).

[244] E. Drakonaki. ‘Ultrasound elastography for imaging tendons and muscles’. In: *Journal of ultrasonography* 12.49 (2012), p. 214.

[245] Q. Zhang, N. C. Adam, S. Hosseini Nasab, W. R. Taylor and C. R. Smith. ‘Techniques for in vivo measurement of ligament and tendon strain: a review’. In: *Annals of biomedical engineering* 49.1 (2021), pp. 7–28.

[246] A. A. Nikooyan, H. Veeger, P. Westerhoff, F. Graichen, G. Bergmann and F. Van der Helm. ‘Validation of the Delft Shoulder and Elbow Model using in-vivo glenohumeral joint contact forces’. In: *Journal of biomechanics* 43.15 (2010), pp. 3007–3014.

[247] M. Glover, E. Montague, A. Pollitt, S. Guthrie, S. Hanney, M. Buxton and J. Grant. ‘Estimating the returns to United Kingdom publicly funded musculoskeletal disease research in terms of net value of improved health outcomes’. In: *Health research policy and systems* 16.1 (2018), p. 1.

[248] M. Pinelli, S. Manetti and E. Lettieri. ‘Assessing the social and environmental impact of healthcare technologies: Towards an extended social return on investment’. In: *International Journal of Environmental Research and Public Health* 20.6 (2023), p. 5224.

[249] M. Kelly, R. L. Donovan, Z. H. Dailiana, H.-C. Pape, F. Chana-Rodríguez, C. Pari, K. J. Ponsen, S. Cattaneo, A. Belluati, A. Contini *et al.* ‘Rehabilitation after musculoskeletal injury: European perspective’. In: *OTA International* 7.5S (2024), e330.

Acknowledgements

At the end of a life phase, it is always good to pause for a moment and reflect on it calmly. Now that my time at TU Delft is coming to an end, it is my turn to do so. I have grown as a person during my PhD journey and have had the amazing opportunity to share it with some of the best travel mates I could have asked for. Honestly, I would choose to come to Delft again and again without a single doubt if I could go back. I have had a great time at CoR and BME, so I need to say some “thank-yous” before leaving, even if they will not nearly be enough.

First, I would like to acknowledge that I would not be here today without Luka and Ajay: you were a constant source of inspiration and support with your profound knowledge, always open to spare some time to get me started in the beginning and to discuss new ideas and next directions later on. Thinking about it now, we have envisaged, discussed, or only drafted so many ideas together that 4 years extra would not be enough!

Ajay, thank you for selecting me as your first PhD student in Delft, and for making me passionate about biomechanics. Your dedication to science and attention to detail have been a real example in these years. You helped me question my results and get to the bottom of everything, and I hope I will keep up your tireless, contagious enthusiasm. Thank you also for suggesting all the group activities, which allowed us to discuss more than science and work. We never went on that bike ride together, maybe one day we will!

Luka, thank you for your ability to lead my research, focusing on what really mattered and generously contributing to it. I can remember many moments in which we worked together on a figure for a paper, or on a clear analogy to explain our research approach - I saw your selflessness and dedication to help your students, and I hope I will be able to match that in the future, wherever I will end up. Thank you for regularly stopping by my desk when deadlines were close (sometimes bringing Slovenian candies!), and for showing with your example that often science requires a fresh mind more than numbing long hours.

Micah, I saw you move up the ranks since I arrived, yet you remained always available and involved in my project, even as more duties piled up. I bet you did not expect that the arm brace you designed would have ended up on the cover of a PhD thesis! Thank you for sharing your knowledge when there was something to debug, and for always contributing to positive discussions and meetings.

David, thank you for finding the time to meet in the moments of need, even if you were always fighting deadlines and winning grants! I admired your ability to be a leader and a passionate researcher at the same time, combining your extraordinary experience to deliver both meaningful ”bird-eye-views” and engage in precise discussions. Discussions with you were always energizing, and you made sure

to always give me an environment where I could thrive.

Naturally, I would also like to thank the other members of my committee, **Prof. van der Helm**, **Prof. Sartori**, **Dr. Marchal Crespo**, and **Dr. Wolf**, for accepting to be part of my PhD defence, and for their interest and feedback on my work. My gratitude goes also to **Prof. de Jong**, for chairing my defense, and to **Prof. Riener** for hosting me at ETH Zürich in his stimulating lab.

I would like to thank also the researchers I collaborated with, specifically **Ibrahim** and **Lisa**: even if we worked together remotely, we had inspiring discussions and shared interesting ideas (and you also helped me fix some bugs in my code as first users!).

Thank you also to the students I had the honour to work with: **Stephan**, **Irene**, **Noah**, **Jevon**, **Guiomar**, **Ted**, **Florian**, **Hiske**, and **Siyang**. All of you made my daily work more enjoyable, and motivated me with your projects, questions, and dreams.

Martijn, Laura M.C., Chris, Cosimo, Carlos, Laura F., Arkady, thank you all for bringing your enthusiasm to me on many occasions. From the CoR Thursday's drinks to casual chats in the corridors, you were always approachable and a source of positive energy in our department.

To all the members of the **PTbot** project, and those with whom I shared the **CBL** meetings and activities, thank you for the fruitful discussions, and for seeking my advices. And to **Adam**: thank you for advising me in numerous occasions, and for helping me improving my code so many times!

Thank you also to **Hanneke**, **Noortje**, **Loulou**, and **Hellen** for helping with administrative matters, en bedankt dat ik mijn Nederlands mocht oefenen.

Then, how to continue? There are many with whom I have shared beautiful moments, so let's start with the people I saw every morning and evening: my great flatmates, **Tomás** and **Anton**. Thank you, guys, for sharing these years with a smile and for helping each other to remember that, even though we lived so close to the university, our lives should be more than that. That we also needed hikes, movies, dinners, saunas, 7:30-AM gym (that became evening gym at a certain point... At least for me!), that squared plates are not that handy, that sometimes the floor needs watering, and that you can always eat more proteins. Of course there is much more, but let's keep it short!

Lorenzo, thank you for being the first one to reach out on my first day, and for inspiring me with your welcoming spirit. You were a constant presence in these years, even when you moved upstairs and downstairs and everywhere with your little cars, I could always come by for a chat, and to organize the next dinner, board game, or body power together. Thank you for your character (onf), and for sharing your opinion when it counted.

Yuxuan, thank you for bringing new energies and enthusiasm to my last year in Delft. It was a lot of fun to see you getting started - I hope that I was not too tedious of a colleague when I pulled you into my projects. I really enjoyed exchanging ideas and working together!

To my Italian flatmate in the flatlands, **Giovanni**, thank you for welcoming me to the most photographed house in the whole Netherlands and for making sure it did not end up in flames. And for the bike rides with the wet wind opposing us, but we always prevailed. You, instead, did not prevail in convincing me to give the GPs a serious try, so

I need to apologize for that! **Beatriz**, thank you for sharing with me your tapas-based dinners, and for bringing reason in the house!

Gustavo, thank you for the motivating breaks in the late evenings, the back-flips, and for introducing me to the Brasilian bbq cult. **Rodrigo**, thanks for sharing your expertise with the Kuka, and for making me wear a tuxedo for the first time. **Julian**, thanks for becoming the unofficial CoR baker, everywhere I go I tell about how lucky we were to have you there. **Manauel**, thanks for encouraging me to always join talks and seminars, and keep a positive, open mind. **Nicky**, dank je wel voor de leuke gesprekken over robotica en het leven, en voor de snelle Nederlandse lunches. **Mariano**, thanks for the refreshing geopolitical insights, fruit breaks, and frisbee throws. **Lasse**, thanks for organizing the (very first?) CoR hackaton, and for all your research suggestions. **Jelle**, thank you for the chats, the uprisings, and for the ice skating. **Taso**, thank you for keeping me company in the crow's nest, where we discovered ChatGPT for the first time. Thanks to all of my office mates over these years: **Ashwin**, **Dennis**, **Max**, **Katie**, **Roel**, **Khaldon**, **Tom**, **Olger**, **Vishal**, **Zhaochong**, and **Pablo**, and thanks **Zlatan** for leaving me your the desk and plants. Thanks to the HRI team, for sharing nice demos and fun moments: **Stefano**, **Renchi**, **Giuseppe**, **Salvo**, **Alex** V.d.B, **Amy**, **Alberto**, and **Jagan**. Thanks **Kate** for your tirelss role at the PhD council. Thanks **Alex** R. for the nighltly stop-overs. Thank you **Beatrice** for sharing with me your exoskeleton wisdom.

To my friends and temporary colleagues at ETH Zürich, thank you all for the warm welcome and the great time during my research visit: especially to **Florence** and **Max-Ole**, the ANYexo gossipers, but also to all the others, for getting me up to speed with your robots, Switzerland, and foosball **Alfie**, **Jan**, **Sara**, **Andri**, **Adrian**, **Michi**, **David**, and all the others!

To the squash group, **Alvaro**, **Natalia**, **Elia**, and **Corrado**: I had probably lowest attendance score, but I definitely enjoyed those games.

To **Ema**, **Erik**, **Menno**, **Dario**, **Mees**, and **Giacomo**: thank you for the wins and the tactics, and for being my trusted blockers on sunny, super hot days, but also in rain and wind where a volleyball should not be thrown. Also thanks to all of my team mates indoor and on the sand at Punch and Shot for the planks, the drills, and the adrenaline!

Sabrina, **Hassan**, **Zoe**, **Judith**, and **Tom**: thank you for running meetings together, thinking on the fly, and discussing common problems. I always had a nice time with you, and I am sorry we ran out of time so quickly...

To **Luzia**, **Anna**, **Saray**, **Maxi**, **Zhaoting**, **Daniel**, **Forough**, **Chuhan**, and **Ebrahim**: thank you for the chill lunch-breaks, and the fun and more serious discussions we had there, and for always sharing a smile. It meant a lot to me.

To **Sara**, thank you for the nice times, in Delft and on the mountains alike!

And to the other members of the RoboCasa, past and present: **Ámbar** and **Vito**, thanks for contributing to the nice atmosphere in the house!

Poi un ringraziamento va anche al leggendario gruppo Risiko, evolutosi e radicatosi nei territori collinari del Varesotto con aspirazioni sudamericane. Grazie a **Edoardo**, **Pietro**, **Veronica**, **Sabrina**, **Stefano**, **Andrea**, **Caterina**, **Chiara** e **Cristina** per aver inaugurato così tanti anni insieme, e per gli incoraggiamenti durante il percorso.

Al miglior cheering team al mondo, il mio ringraziamento più grande: siamo capaci

di rimanere uniti nelle gioie e nelle difficoltà, di ascoltarci con attenzione e di capirci al volo. **Sabina**, grazie per avermi sostenuto nell'inseguire il mio futuro ed i miei sogni, scoccandomi lontano da te nel momento più difficile. Grazie perché mi hai donato la vita, forza e curiosità, e sostegno in ogni momento. **Marina**, grazie perchè siamo vicini anche se lontani: cittadini del mondo, sapremo sempre ritrovarci. **Elena**, grazie per le telefonate e i consigli, e per la premura con cui ti sei sempre presa cura di me. **Vanni**, grazie che ci sei stato, anche se solo all'inizio di questo viaggio: ora lassù sei stella forte che mi illumina.

I por fin, **Maia**, it is no secret that, without you, I would have not stood where I stand today. You were the reason I came back to Delft, and thanks to you I lived these years to the fullest. Looking back to our footprints together, we really walked a lot since our beginning, and every step has been marvelous and unique. Thank you for being my wife and muse, ready to wear every kind of boots to follow me in my adventures. I will always do the same for you.

I want to close with a reflection: I thought hard, but I do not have any **anti-acknowledgements** to make. However, I am aware that this is probably because I am a European male, and the fact that did not experience discrimination does not mean that everyone around me feels the same. I strongly believe that diversity and inclusion are essential to make science and society thrive, and would like to thank everyone who has opened my mind on how important it is to always see things from different perspectives. May we strive to look at our neighbors as God looks at them, and not use our own limited experience to judge them.

Italo Belli, Basel, January 2026th

Curriculum Vitæ

Italo Belli

23-02-1996 Born in Varese, Italy.

Education

2011–2015 Scientific Highschool ”Galileo Ferraris”
Varese, Italy

2015–2018 Bachelor of Science & Automation Engineering
Politecnico di Milano, Italy

2018–2020 Masters of Science & Automation and Control Engineering
Politecnico di Milano, Italy
Thesis: Optimization-based hybrid wheeled-legged locomotion for a quadrupedal robot
Supervisors: Prof. dr. P. Rocco, dr. M. Parigi Polverini

2019 (Feb.-July) Erasmus Exchange Student
Technische Universiteit Delft, The Netherlands

2021–2025 PhD. in Robotics
Technische Universiteit Delft, The Netherlands
Thesis: Biomechanics-aware control for robot-assisted physiotherapy
Promotor: Prof. dr. ir. D.A. Abbink
Copromotors: Prof. dr. ir. L. Peternel, Prof. dr. ir. A. Seth

Experience

2020 Intern, Humanoid and Human-Centred Mechatronics (IIT, Italy)

2020–2021 Research Fellow, Artificial and Mechanical Intelligence (IIT, Italy)
2025 (July-Dec.) Visiting Researcher, Sensory-Motor Systems Lab (ETH Zurich, Switzerland)

List of Publications

Journal Papers

- L. Noteboom, I. Belli, M. J. Hoozemans, A. Seth, H. Veeger and F. Van Der Helm. 'Effects of bench press technique variations on musculoskeletal shoulder loads and potential injury risk'. *Frontiers in Physiology*, 2024
- 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*, 2023 – **Chapter 2**

Conference Proceedings

- I. Belli, Y. Hu, J. M. Prendergast, D. Abbink, A. Seth and L. Peternel. 'A Shared Control Approach to Safely Limiting Patient Motion Based on Tendon Strain During Robotic-Assisted Shoulder Rehabilitation'. *2025 International Conference On Rehabilitation Robotics (ICORR)*, 2025 – **Chapter 5**
- J. Ravenberg, I. Belli, J. M. Prendergast, A. Seth and L. Peternel. 'Creating Discomfort Maps via Hand-held Human Feedback Interface for Robotic Shoulder Physiotherapy'. *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024 – **Chapter 6**
- 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'. *2023 IEEE-RAS 22nd International Conference on Humanoid Robots (Humanoids)*, 2023 – **Chapter 4**
- 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'. *2022 IEEE-RAS 21st International Conference on Humanoid Robots (Humanoids)*, 2022

Pre-prints

- I. M. I. Hasan, I. Belli, A. Seth and E. M. Gutierrez-Farewik. 'Modeling glenohumeral stability in musculoskeletal simulations: A validation study with in vivo contact forces'. *bioRxiv*, 2025
- I. Belli, J. M. Prendergast, A. Seth and L. Peternel. 'Biomechanics-Aware Trajectory Optimization for Navigation during Robotic Physiotherapy'. *arXiv preprint arXiv:2411.03873*, 2024 – **Chapter 3**