

UNIVERSITÀ
DI PAVIA

Faculty of Engineering

Department of Industrial and Information Engineering

Master's Thesis in Industrial Automation Engineering – Industrial Technologies and Management

Human Joint Trajectory Tracking with Patient-Specific Range-of-Motion Constraints for Robotic Rehabilitation

Candidate:
Luca Attilio Nascimbene

Supervisor:
Prof. Antonella Ferrara

Co-Supervisor:
PhD. Chiara Alessi

A.Y. 2025/2026

Abstract

End-effector-based robotic systems represent a promising solution for upper-limb rehabilitation due to their mechanical simplicity, adaptability and ease of use. Since the interaction with the patient occurs only at the end-effector, however, reproducing a desired Cartesian trajectory does not guarantee that the patient's joints evolve as intended. The Human Joint Trajectory Tracking (HJTT) framework addresses this limitation by generating robot trajectories that minimize the discrepancy between a reference joint configuration and the one induced in the patient.

The reference motion is recorded from a therapist whose mobility is unimpaired. When the patient's Range of Motion (ROM) is restricted, reproducing it faithfully may therefore require joint excursions beyond those that the patient can safely achieve, with possible discomfort or injury. The objective of this thesis is to adapt the trajectory to the ROM of the individual patient.

The adaptation is formulated as an optimization problem in the human joint space, in which the original tracking objective is combined with logarithmic barrier terms associated with the ROM boundaries. The formulation supports simultaneous constraints on several joints and accounts for differences between their admissible ranges. Reference trajectories are acquired with the Xsens MVN Awinda inertial motion capture system, and the adapted motions are executed on a Franka Emika Panda collaborative robot.

For each reference trajectory, the algorithm computes the robot joint trajectory while simulating the corresponding evolution of the human joint configuration. The predicted human joint trajectories remained within the prescribed ROM limits throughout the adapted motions. The resulting trajectories were then executed by the robot with a single healthy subject under imposed ROM limits. The experimental measurements showed the intended reduction in joint excursion on the constrained joints, with a corresponding deviation from the reference.

The resulting framework extends HJTT toward safer and more patient-specific end-effector-based rehabilitation, while retaining the joint-space tracking formulation of the original method.

Sommario

I sistemi robotici basati su end effector rappresentano una soluzione promettente per la riabilitazione dell'arto superiore grazie alla loro semplicità meccanica, adattabilità e facilità d'uso. Tuttavia, poiché l'interazione con il paziente avviene esclusivamente a livello dell'end effector, riprodurre una traiettoria cartesiana desiderata non garantisce che le articolazioni del paziente si muovano come previsto. L'algoritmo Human Joint Trajectory Tracking (HJTT) supera questa limitazione generando traiettorie del robot che riducono al minimo la differenza tra la configurazione articolare di riferimento e quella indotta nel paziente.

Il movimento di riferimento viene registrato da un terapista che non ha limitazioni di mobilità. Quando il range di movimento (ROM) del paziente è ridotto, riprodurlo fedelmente può richiedere escursioni articolari superiori a quelle che il paziente può raggiungere in sicurezza, causando possibile disagio o rischio di lesione.

L'obiettivo di questa tesi è quindi quello di adattare la traiettoria al ROM del singolo paziente.

Tale adattamento è formulato come un problema di ottimizzazione nello spazio articolare umano, in cui l'obiettivo di inseguimento originale è combinato con termini di barriera logaritmica associati ai limiti del ROM. Questa formulazione permette di vincolare più giunti simultaneamente e di considerare le differenze tra i rispettivi intervalli ammissibili. Le traiettorie di riferimento sono acquisite tramite il sistema di motion capture inerziale Xsens MVN Awinda, mentre i movimenti adattati vengono eseguiti su un robot collaborativo Franka Emika Panda.

Per ciascuna traiettoria di riferimento, l'algoritmo calcola la corrispondente traiettoria articolare del robot, simulando contestualmente l'evoluzione della configurazione articolare umana.

Le traiettorie articolari umane predette sono rimaste entro i limiti del ROM prescritti per tutta la durata dei movimenti adattati. Successivamente, le traiettorie ottenute sono state eseguite dal robot su un singolo soggetto sano con limiti di ROM imposti. Le misure sperimentali hanno mostrato la riduzione attesa dell'escursione articolare sui giunti vincolati, con una corrispondente deviazione rispetto al riferimento.

Il metodo risultante estende l'HJTT verso una riabilitazione basata su end effector più sicura e specifica per il paziente, mantenendo la formulazione di inseguimento nello spazio articolare del metodo originale.

Declaration on the Use of Generative Artificial Intelligence

In accordance with the guidelines of the University of Pavia, I declare that I used generative artificial intelligence tools, specifically ChatGPT (OpenAI) and Claude (Anthropic), in an assistive capacity during the preparation of this thesis.

These tools were used to support the following activities:

- linguistic revision and editing of the text, including the rephrasing of selected passages based on the author's notes, source code, and simulation results;
- conversion of the document from Microsoft Word to LaTeX;
- review and debugging of the source code;
- processing and organization of simulation output files.

All AI-assisted outputs were reviewed and, where necessary, corrected and validated by the author. The scientific objectives, methodology, interpretation of results, selection of sources, and conclusions are the author's own work. The author retains full responsibility for the content of this thesis.

Contents

1 Introduction **1**
- 1.1 Robotic rehabilitation of the upper limb . . . 1
- 1.2 Robotic devices for rehabilitation . . . 1
  - 1.2.1 Exoskeletons . . . 1
  - 1.2.2 End-effector devices . . . 2
- 1.3 Trajectory tracking in end-effector rehabilitation . . . 2
  - 1.3.1 Cartesian Teach and Track . . . 3
  - 1.3.2 Human Joint Trajectory Tracking . . . 3
- 1.4 Range of motion limitations in patients . . . 4
- 1.5 Thesis contribution . . . 4
- 1.6 Thesis structure . . . 6

2 Fundamentals of Robotics and Optimization **7**
- 2.1 Basics of Robotics . . . 7
- 2.2 Manipulator kinematics . . . 8
  - 2.2.1 Direct Kinematics . . . 8
  - 2.2.2 Differential Kinematics . . . 10
  - 2.2.3 Inverse Kinematics . . . 12
- 2.3 Redundancy and damped pseudoinverse . . . 13
  - 2.3.1 Kinematic redundancy . . . 13
  - 2.3.2 Pseudoinverse solutions and the null space . . . 14
  - 2.3.3 Kinematic singularities and damped least squares . . . 15
- 2.4 Numerical optimization: Newton, Gauss–Newton and barrier methods 15
  - 2.4.1 Unconstrained minimization and Newton's method . . . 16
  - 2.4.2 Gauss–Newton methods for nonlinear least squares . . . 17
  - 2.4.3 Barrier methods for inequality constraints . . . 18

3 Materials **19**
- 3.1 Franka Emika Panda . . . 19
- 3.2 Xsens MVN Awinda . . . 22
  - 3.2.1 Coordinate systems and joint angle conventions . . . 23

4 Problem Statement and Modelling **27**
- 4.1 Human arm model . . . 27
- 4.2 Baseline Human Joint Trajectory Tracking Algorithm . . . 29
  - 4.2.1 Kinematic coupling between the human arm and the robot . . . 30
  - 4.2.2 Human joint tracking problem . . . 31

4.2.3 Iterative robot update 32
4.2.4 Update of the human configuration 33
4.2.5 Application to a complete trajectory 34
4.3 Human Joint Range of Motion and Patient-Specific Limits 35
4.4 Problem Statement: Enforcing Patient-Specific ROM Constraints . . 36

5 ROM-Constrained HJTT: Proposed Method 39
5.1 Cost Function with Logarithmic Barrier 39
5.2 Gradient and Hessian in the Human Joint Space 40
5.3 Local HJTT Mapping and Generalized Gauss–Newton Approximation 41
5.4 Gauss–Newton Iterative Update and Robotic Realization 42
5.5 Multiple Constrained Joints and ROM Normalization 45
5.6 Feasibility, Stopping Criterion and Relation with the Baseline HJTT . 45
5.7 Application to a complete trajectory 48

6 Experimental Setup 49
6.1 Overview and objectives 49
6.1.1 Subject description 49
6.1.2 Acquisition and reference trajectories 50
6.2 Definition of the range-of-motion limits 51
6.3 Algorithm parameters and numerical safeguards 52
6.4 Experimental implementation 53
6.5 Software integration and experimental workflow 55

7 Results 57
7.1 Introduction to results 57
7.2 Predicted human configuration 57
7.3 Experimental results 61
7.3.1 Movement 1 61
7.3.2 Movement 2 63
7.4 Overall observations 64

8 Conclusion 65
8.1 Conclusions 65
8.2 Future Work 66

List of Figures

1.1 Graphical abstract of the thesis. Above, the inputs and the output of the ROM-constrained HJTT formulation; below, the experimental workflow adopted for its validation. 5

3.1 Kinematic chain of the Franka Emika Panda with the frames assigned according to the modified Denavit–Hartenberg convention, in the configuration $\theta_1 = \cdots = \theta_7 = 0$. Image from the Franka Control Interface documentation [1]. 21
3.2 Flange of the Franka Emika Panda with the custom 3D-printed handle employed in the experimental setup. The cylindrical grip is 150 mm long and has a radius of 20 mm. 21
3.3 MVN Awinda system: full-body configuration worn by a subject (left) and the set of wireless MTw trackers (right). Images courtesy of Xsens Technologies B.V. [2]. 22
3.4 Selection of the calibration pose in MVN Analyze. The N-pose, on the left, is the pose recommended by the manufacturer and the one used in this work. 23
3.5 Orientation of the anatomical body frame at the right shoulder in the neutral pose. Rotations about the z-, x- and y-axis correspond respectively to flexion–extension, abduction–adduction and internal–external rotation. Adapted from [3]. 24

4.1 Kinematic model of the human upper limb: the rotation axes of the three shoulder joints and of the elbow, the segment lengths l_{UA} and l_{FA}, and the frame attached to the hand. Reproduced from [4]. 29
4.2 Block diagram of the baseline HJTT framework: the joint trajectory $q_h^\star$ recorded from the therapist and the Denavit–Hartenberg parameters of the patient are combined to generate the robot joint trajectory $q_r^\star$ executed during the exercise. Adapted from [5]. 30

5.1 Iteration of the ROM-constrained formulation for a single sample of the reference trajectory. Blocks in blue are those introduced in this work; the remaining ones belong to the baseline HJTT algorithm. . . 44

6.1 Joint trajectories of the two reference movements, recorded from the therapist in free space. Left: Movement 1. Right: Movement 2. The excursions of these trajectories are the ones reported in Table 6.2. . . 51

6.2 Experimental setup in the Intelligent Robotics Laboratory. The robot base is 0.90 m above the floor and the seat of the subject is placed at 0.90 m from it. 53
6.3 Initial, intermediate and final instants of the guided execution of the *constrained* trajectory, for the two movements. Top row: Movement 1. Bottom row: Movement 2. The intermediate frame of each row corresponds to the peak of the dominant constrained joint, the shoulder flexion–extension in the first movement and the abduction–adduction in the second. 54
6.4 Schematic representation of the experimental workflow. 55

7.1 Movement 1: human joint configuration predicted during trajectory generation (solid green), superimposed on the therapist reference (solid blue). Dashed lines indicate the imposed ROM limits on the constrained joints. 58
7.2 Movement 2: predicted human joint configuration, with the same representation adopted in Figure 7.1. 59
7.3 End-effector trajectory generated with and without the ROM constraint, for the two movements. The three axes are drawn to the same scale. 61
7.4 Movement 1: measured human joint trajectories. Mean over the ten constrained repetitions (solid green), min–max interval across trials (shaded), therapist reference and unconstrained execution. Dashed lines indicate the imposed ROM limits on the constrained joints. . . . 62
7.5 Movement 2: measured human joint trajectories, with the same representation adopted in Figure 7.4. 63

List of Tables

3.1 Joint limits and kinematic constraints of the Franka Emika Panda [1]. 19
3.2 Denavit–Hartenberg parameters adopted for the nominal Panda kinematic model [4]. 20

4.1 Denavit–Hartenberg parameters of the human arm model, expressed in the standard convention [5, 4]. 28

6.1 Anthropometric parameters of the subject. The shoulder height is measured from the floor with the subject seated, in the fixed chair position adopted throughout the experiments. 50
6.2 Joint excursions of the two reference movements, measured on the recorded trajectories. 50
6.3 Range-of-motion limits imposed in the two movements, obtained from (6.1) with $\rho_{\text{low}} = 0$ and $\rho_{\text{high}} = 0.80$. All values in degrees. 51
6.4 Parameters of the constrained algorithm. 52

7.1 Movement 1: peak values of the constrained joints [deg]. Angles are reported according to the anatomical sign convention adopted in the experimental plots. 62
7.2 Movement 2: peak values of the constrained joints [deg]. The SAA values are reported with the anatomical sign, obtained from the kinematic coordinate through the sign transformation of Section 4.1. . . . 63

List of Acronyms

ADL	Activities of Daily Living
CTT	Cartesian Teach and Track
DH	Denavit–Hartenberg
DK	Differential Kinematics
DLS	Damped Least Squares
DoF	Degree of Freedom
EFE	Elbow Flexion/Extension
FCI	Franka Control Interface
FER	Franka Emika Robot
FK	Forward Kinematics
GGN	Generalized Gauss–Newton
GUI	Graphical User Interface
HJTT	Human Joint Trajectory Tracking
IK	Inverse Kinematics
IMU	Inertial Measurement Unit
ISB	International Society of Biomechanics
MTw	Wireless Motion Tracker
NWU	North-West-Up
ROM	Range of Motion
ROS	Robot Operating System
SAA	Shoulder Abduction/Adduction
SDI	Strap Down Integration
SFE	Shoulder Flexion/Extension
SIER	Shoulder Internal/External Rotation
TT	Teach and Track

Chapter 1

Introduction

1.1 Robotic rehabilitation of the upper limb

Motor impairment of the upper limb is among the most frequent and persistent consequences of stroke, and its recovery depends on practice that is repeated many times, delivered at sufficient intensity and directed at meaningful tasks. Robotic devices are attractive in this context because they can administer precisely this kind of practice, sustaining a high number of repetitions and a consistent level of training throughout extended sessions [6].

Beyond the delivery of movement, therapeutic robots are motivated by two objectives: sustaining recovery through therapy that can be dosed and controlled, and providing new means of evaluating the progress of the patient [7]. A third motivation is organizational: by sustaining the repetitive component of training, robotic devices can reduce the therapist's workload and facilitate the delivery of intensive treatment [8]. Compared with conventional therapy, a robotic system can reproduce the prescribed exercise at a controlled intensity, graduate the assistance it provides according to the extent of the impairment, and record during the session the quantitative information the clinician needs in order to monitor the patient and adjust the treatment [8].

It should be noted that the clinical evidence on robotic therapy is not yet conclusive, and the optimal intensity of treatment remains an open question [6]. The interest of the field therefore lies less in replacing conventional therapy than in complementing it, by making intensive practice sustainable and its outcome measurable.

1.2 Robotic devices for rehabilitation

Rehabilitation robots are conventionally classified into two families according to their mechanical structure and to the way they interact with the body: exoskeletons and end-effector devices [9, 10].

1.2.1 Exoskeletons

Exoskeletons are wearable systems that reproduce the kinematic structure of the human limb. Their joints are actuated and aligned with the anatomical axes of the

patient, so that each articulation can be commanded individually. This makes them well suited to targeted rehabilitation, in which the motion of a specific joint must be isolated and trained, and it allows the prescribed motion to be expressed directly in the human joint space.

The price of this capability is mechanical: exoskeletons are complex, heavy, and difficult to align accurately with the anatomy of an individual patient. Misalignment reduces comfort, limits adaptability across body types and lengthens the setup time before each session.

1.2.2 End-effector devices

End-effector devices interact with the patient at a single contact point — typically a handle — while the motion of the whole limb results from the mechanical coupling between the hand and the rest of the arm. They are mechanically simpler and faster to prepare and, because the interaction takes place at a single interface, they generally require less mechanical adaptation to the anatomy of the individual patient than exoskeletons [8]. A well-established example in upper-limb rehabilitation is the MIT-MANUS [7].

The counterpart of this simplicity is that the individual joints of the patient are neither commanded nor measured directly. The robot controls where the hand is; how the arm arranges itself to place the hand there is left to the patient's anatomy. This distinction is the starting point of the present work.

A further consideration concerns the hardware itself. Devices designed specifically for rehabilitation require dedicated mechanical structures, which makes them costly and difficult to deploy outside research laboratories. Collaborative industrial manipulators offer an alternative: they are designed to share a workspace with human operators, are commercially available and are supported by standard software frameworks [11, 12, 13]. The versatility of collaborative manipulators has also been demonstrated in physical human–robot interaction tasks, such as ergonomic object handover, where reference generation and adaptive control strategies have been experimentally evaluated on a Franka Emika Panda robot [14]. Recent work has shown that such manipulators can be integrated into a clinician-oriented rehabilitation workflow, with an interface that allows therapists to record, analyze and design patient-specific exercises without programming expertise [15]. The same line of work has since been consolidated into a single platform that brings assessment, force estimation and trajectory tracking together under one clinician-oriented interface, and that reports the operational barriers emerging when such a system is deployed in a real clinical environment [16].

1.3 Trajectory tracking in end-effector rehabilitation

Robotic rehabilitation can be performed under different interaction modalities. Depending on the residual capabilities of the patient, the robot may oppose the movement, assist it as required, or impose it entirely by guiding the limb along a pre-

scribed trajectory [8]. This work focuses on the latter case, namely passive rehabilitation, in which voluntary contribution is absent or limited and the prescribed robot motion plays a dominant role in determining the movement of the limb.

Teach and Track (TT) is a procedure commonly proposed in rehabilitation settings. The motion to be administered is not programmed but obtained by demonstration: the therapist physically guides the robot along the desired movement, which is recorded and subsequently reproduced on the patient. The appeal of the approach is that the exercise is defined by the clinician through the movement itself, without requiring any programming. This work is concerned with this family of procedures. What is recorded, and what is subsequently enforced, determines the therapeutic content of the exercise.

1.3.1 Cartesian Teach and Track

In the established Teach and Track the pose of the end-effector is recorded during the demonstration and reproduced during the execution, solving the inverse kinematics of the manipulator to obtain the corresponding joint references. This strategy, referred to as Cartesian Teach and Track (CTT), aims to reproduce the demonstrated end-effector trajectory with high accuracy in the Cartesian space [5, 4].

It does not, however, guarantee the resulting human joint motion. The same Cartesian path can be realized by markedly different arm configurations, and which one occurs depends on the morphology of the person holding the handle. When therapist and patient differ in segment lengths, reproducing the recorded path faithfully may therefore induce a joint motion that differs from the demonstrated one, with consequences for safety, comfort and therapeutic relevance.

This limitation does not arise in the same form with exoskeletons, where the prescribed motion can be expressed directly in the human joint space and individual joints can be controlled according to the recorded trajectory [17, 18, 19]. The difficulty is specific to end-effector systems, and it is the consequence of interacting with the limb at a single point.

1.3.2 Human Joint Trajectory Tracking

The Human Joint Trajectory Tracking (HJTT) framework addresses this limitation without abandoning the end-effector architecture [5]. Instead of recording the Cartesian pose of the hand, the joint angles of the demonstrator are acquired through inertial motion capture and used as the reference. During execution, the robot configuration is computed so as to minimize the discrepancy between this reference and the joint configuration induced in the patient, exploiting a parametric kinematic model of the human arm to account for the individual anthropometry.

The quantity that is minimized is therefore the human joint error rather than the pose error of the end-effector, and the adaptation to the patient occurs through the kinematic model rather than through the mechanical structure of the device. This framework, which has been experimentally evaluated on subjects with different arm morphologies [4], constitutes the starting point of the present work.

1.4 Range of motion limitations in patients

The reference motion is demonstrated by a therapist, that is, by a person whose mobility is unimpaired. The patient who receives the therapy is, in general, not in that condition. Limited Range of Motion (ROM) is a frequent consequence of stroke, produced by the combined effect of spasticity, loss of muscular strength and the development of joint contractures, in proportions that differ from one patient to another; in hemiplegia the two sides of the body are affected to a different degree [20].

The functional relevance of these limitations has been quantified in the literature. A systematic review of the joint excursions that unimpaired participants employ while carrying out activities of daily living places the requirement at roughly 130° at the shoulder, in both flexion and abduction, while several everyday tasks such as feeding and personal hygiene call for the elbow to be flexed close to the end of its range, at about 150° [21]. Restrictions in these ranges can therefore interfere with activities of daily living and reduce the functional independence of the patient, which is why the recovery of range of motion represents an important rehabilitation objective.

This introduces a difficulty that is specific to a framework designed to reproduce joint trajectories faithfully. The more accurately the algorithm tracks the demonstrated joint motion, the more closely it may reproduce excursions that exceed the patient's available range of motion. Accuracy of tracking and safety of the resulting motion are, in this situation, in tension.

The point is particularly relevant in passive rehabilitation, where the robot is required to impose the prescribed motion: control laws developed for this purpose may treat the force exerted by the patient as a disturbance to be compensated rather than as information about the state of the limb [19]. Sliding-mode observation and control schemes developed for the same manipulator and the same exercise are designed precisely to estimate that interaction force and reject it during passive execution [22], which is the appropriate objective at the control level but leaves the admissibility of the prescribed motion untouched. Compatibility with the available range of motion should therefore be accounted for when the trajectory is generated, rather than left to the behavior of the low-level controller.

Safety and individualization have been addressed in the literature, although by different means. A representative approach defines an admissible region statistically, learning it from the motion of healthy subjects, and reshapes the trajectory online inside that region while tracking it under a variable impedance law [17]. That formulation is developed for an exoskeleton, where the joints are commanded directly, and the admissible region is learned statistically from demonstrated motion rather than defined from patient-specific anatomical ROM bounds.

1.5 Thesis contribution

This thesis extends the Human Joint Trajectory Tracking framework to account for patient-specific range-of-motion limits during trajectory generation. The adaptation is formulated as an optimization problem in the human joint space, in which the

tracking objective of the original method is combined with logarithmic barrier terms associated with the boundaries of the admissible range. The resulting problem is solved by a generalized Gauss–Newton scheme, and the correction defined in the human joint space is mapped to the robot through the differential coupling between the two kinematic chains, preserving the kinematic consistency between the motion of the robot and the human joint motion it induces.

The formulation admits constraints on several joints simultaneously and includes a ROM-dependent scaling of the barrier terms, so that their curvature depends on the relative proximity to the joint limits rather than directly on the absolute width of each admissible range. This scaling therefore supports the use of a common barrier weight across the constrained joints. The method reduces to the original algorithm when the constraints are removed, and therefore extends it rather than replacing it.

The proposed method was applied to reference trajectories acquired with the Xsens MVN Awinda inertial motion capture system and experimentally evaluated on a Franka Emika Panda collaborative manipulator with a healthy subject under artificially imposed ROM limitations.

Figure 1.1 summarizes both the contribution and the way it was evaluated. The upper part shows the inputs of the method, namely the joint trajectory demonstrated by the therapist, the range-of-motion limits prescribed for the patient and the kinematic parameters of the two chains, together with the robot joint trajectory that the method returns. The lower part shows the sequence followed in the experiments: recording of the movement, offline generation of the trajectory and execution on the manipulator.

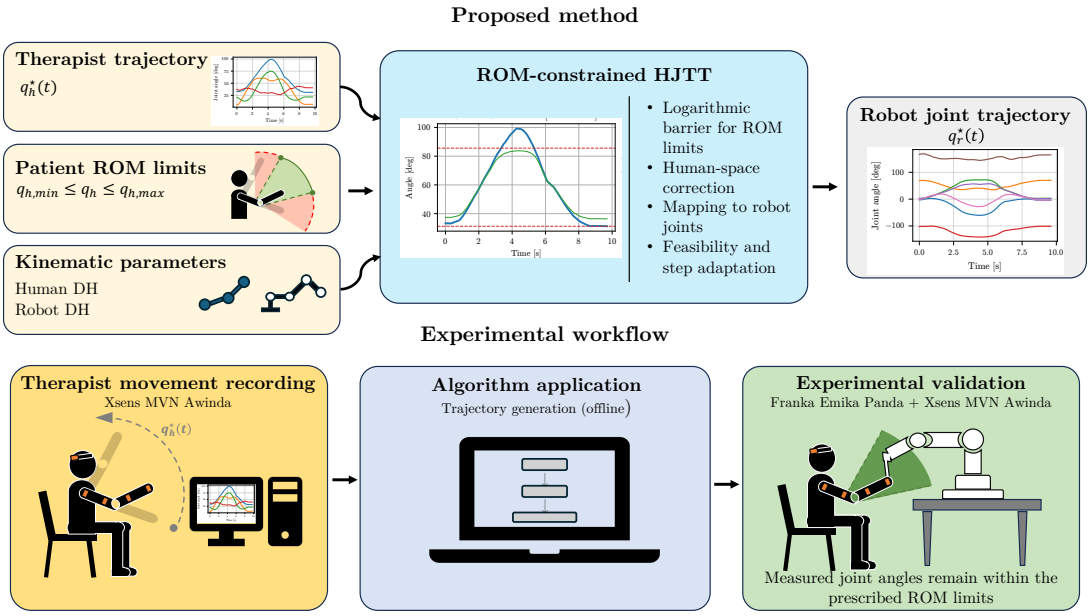


Figure 1.1: Graphical abstract of the thesis. Above, the inputs and the output of the ROM-constrained HJTT formulation; below, the experimental workflow adopted for its validation.

1.6 Thesis structure

This thesis is structured as follows:

- In **Chapter 2** the robotics and optimization concepts used throughout the work are introduced: the direct, differential and inverse kinematics of a manipulator, the treatment of redundant chains through the pseudoinverse and its damped form, and the numerical methods employed in the following chapters, namely the Newton and Gauss–Newton schemes and the barrier approach to inequality constraints.
- In **Chapter 3** the experimental equipment is described: the collaborative manipulator used to execute the trajectories and the inertial motion capture system used to record them, with particular attention to the reference frames and to the joint-angle conventions that determine how the acquired data are interpreted.
- In **Chapter 4** the kinematic model of the human arm is introduced, the baseline Human Joint Trajectory Tracking algorithm is recalled in the form in which it was originally proposed, the patient-specific joint range-of-motion constraints are formalized, and the problem addressed by this thesis is stated together with the requirements that a solution must satisfy.
- In **Chapter 5** the proposed method is developed. The tracking objective is augmented with a logarithmic barrier associated with the joint limits, the resulting problem is solved by a generalized Gauss–Newton scheme, and the correction is realized on the manipulator through the differential coupling adopted by the baseline algorithm.
- In **Chapter 6** the experimental setup is described: the subject, the acquisition of the reference movements, the procedure adopted to define the limits, the parameters of the algorithm and the integration with the robot control software.
- In **Chapter 7** the experimental results are presented, for each of the two movements considered, together with the verification of the compliance with the imposed limits and the comparison between the predicted and the measured joint motion.
- In **Chapter 8** the conclusions of the work are drawn and possible directions for further development are presented.

Chapter 2

Fundamentals of Robotics and Optimization

This chapter introduces the concepts of robotics and numerical optimization that are used throughout this work. These range from describing the motion of a manipulator to solving the optimization problems that arise from it.

First, the fundamentals of robotics are presented, with a focus on the mechanical structure of robot manipulators. Next, the kinematics of a manipulator are introduced in their direct, differential and inverse formulations, relating the pose and motion of the end-effector to the joint variables. This is then extended to include redundant manipulators, for which the Jacobian cannot be inverted, by introducing pseudoinverse solutions and their damped counterparts. Finally, the numerical optimization methods used in subsequent chapters are reviewed: the Newton and Gauss–Newton methods for unconstrained minimization, and the barrier approach for problems subject to inequality constraints.

2.1 Basics of Robotics

In an industrial context, a robot is a machine that can be reprogrammed to perform tasks such as manipulation and handling, able to interact with its surroundings to a certain degree and possessing a certain degree of autonomy. The term robotic system refers to the robot together with all its equipment: the manipulator, that is the mechanical structure, with a fixed or mobile base; the power supply; the controller; the sensors and devices required to perform the task; and the communication interfaces managed by the control system [23].

The mechanical structure of a manipulator consists of a series of rigid elements connected by rotational or translational joints. It is described as serial when all the elements are connected so as to form an open kinematic chain. The first element of the chain, known as the base, is usually fixed, while the last one, the end-effector, performs the desired action.

The main components of the kinematic chain are [23, 24]:

- **Joints:** these provide the mobility of the manipulator. They normally have one degree of freedom and are divided into prismatic joints, which allow a rectilinear translation, and revolute joints, which allow a relative rotation.

- **Links:** the rigid elements that form the structure of the robot, connected to one another by the joints.
- **End-effector:** the terminal element or tool through which the manipulator interacts with the environment and performs the assigned task.

2.2 Manipulator kinematics

Robot kinematics is the analysis of the motion of a robotic manipulator without considering the forces or moments that cause it. Kinematics can be divided into three categories: direct, inverse and differential. The direct kinematics aims to calculate the pose of the end-effector as a function of the joint variables. Inverse kinematics involves determining the joint positions that correspond to a given end-effector position and orientation. In order to study the set of poses reachable by the end-effector, the following spaces are defined [24]:

The joint space is a mathematical space representing the set of values that the joint coordinates can take.

The workspace represents the set of poses that can be achieved by the end-effector.

2.2.1 Direct Kinematics

Direct Kinematics, also referred to as Forward Kinematics (FK), addresses the problem of determining the position and the orientation of a manipulator's end-effector once the values of all joint variables are known. It therefore describes the mapping from the joint space of the mechanism to the operational space in which the task is defined.

Considering an open kinematic chain composed of n joints connecting $n+1$ links, with the first link attached to a fixed base, the corresponding joint variables are collected in the vector

$$q = \begin{bmatrix} q_1 & q_2 & \dots & q_n \end{bmatrix}^\top \in \mathbb{R}^n, \tag{2.1}$$

where each q_i denotes an angular displacement for a revolute joint or a linear displacement for a prismatic joint.

Reference frames and homogeneous transformations

In order to describe the chain, a Cartesian frame is attached to each link, so that the relative placement of two consecutive links reduces to a rigid transformation [24]. Given two frames i and j, the transformation carrying coordinates expressed in the former into coordinates expressed in the latter is written as

$$T_i^j = \begin{bmatrix} R_i^j & p_i^j \\ 0^\top & 1 \end{bmatrix} \in SE(3), \tag{2.2}$$

where $R_i^j \in SO(3)$ describes the relative orientation of the two frames and $p_i^j \in \mathbb{R}^3$ the offset between their origins. The vector of coordinates is augmented with a

fourth unit component, so that rotation and translation are applied through a single matrix product rather than as separate operations.

Because rigid transformations compose by multiplication, the pose of the end-effector frame with respect to the base is obtained by traversing the chain link by link,

$$T_e^b(q) = T_0^b \, T_1^0(q_1) \, T_2^1(q_2) \dots T_n^{n-1}(q_n) \, T_e^n, \tag{2.3}$$

where T_0^b and T_e^n are constant matrices relating the base frame to the first link and the last link to the end-effector. Each of the remaining factors depends on one joint variable only, which is what makes the mapping $q \mapsto T_e^b$ computable in closed form.

Denavit–Hartenberg Convention (D–H)

The factors in (2.3) could in principle be derived case by case from the geometry of the mechanism, but this becomes impractical as the number of links grows. The convention introduced by Denavit and Hartenberg removes the arbitrariness by prescribing how the link frames are to be placed, and by reducing the description of the relative placement of two consecutive frames to four scalar quantities [25, 24].

Denoting by joint i the joint connecting link $i-1$ to link i, the frame attached to link i is placed so that its z-axis lies along the axis of joint $i+1$; consequently, the axis of joint i is z_{i-1}. The x-axis of frame i lies along the common normal to z_{i-1} and z_i, and the y-axis completes a right-handed triad. Once the frames are fixed in this way, four parameters suffice to relate frame i to frame $i-1$:

- the distance a_i measured along the common normal;
- the angle α_i between the two joint axes measured about that normal;
- the displacement d_i along z_{i-1};
- the angle θ_i between the two common normals measured about z_{i-1}.

These are referred to as link length, link twist, link offset and joint angle, respectively.

The transformation between two consecutive frames then takes the standardized form

$$A_i^{i-1}(q_i) = \begin{bmatrix} c_{\theta_i} & -s_{\theta_i} c_{\alpha_i} & s_{\theta_i} s_{\alpha_i} & a_i c_{\theta_i} \\ s_{\theta_i} & c_{\theta_i} c_{\alpha_i} & -c_{\theta_i} s_{\alpha_i} & a_i s_{\theta_i} \\ 0 & s_{\alpha_i} & c_{\alpha_i} & d_i \\ 0 & 0 & 0 & 1 \end{bmatrix}, \tag{2.4}$$

with $c_{(\cdot)}$ and $s_{(\cdot)}$ denoting cosine and sine. For a revolute joint the variable is θ_i and the other three parameters are constant; for a prismatic joint the roles of θ_i and d_i are exchanged.

The convention described here is the standard one, and it is the convention adopted for the kinematic model of the human arm introduced in Chapter 4. The

model of the manipulator employs instead the modified Denavit–Hartenberg convention, in which the frames are placed differently and the resulting elementary transformation takes a different form; the corresponding parameters are reported in Chapter 3.

Substituting (2.4) into (2.3), the direct kinematics of the chain can finally be written in compact form as

$$T_e^b(q) = \begin{bmatrix} R_e^b(q) & p_e^b(q) \\ 0^\top & 1 \end{bmatrix} = T_0^b \left(\prod_{i=1}^{n} A_i^{i-1}(q_i) \right) T_e^n, \tag{2.5}$$

which makes explicit both the block structure of the resulting pose and its dependence on the joint vector.

For convenience, the pose is described as a vector and not as a matrix,

$$x_e^b = \begin{bmatrix} p_e^\top & o_e^\top \end{bmatrix}^\top \in \mathbb{R}^6, \tag{2.6}$$

where $p_e \in \mathbb{R}^3$ is the position of the origin of the end-effector frame and $o_e \in \mathbb{R}^3$ encodes its orientation. Since a rotation matrix carries nine entries subject to orthonormality constraints, a three-parameter description is needed to obtain a minimal representation. A representation commonly adopted in inverse-kinematics computations is the axis–angle form, also known as exponential coordinates, in which a rotation is described by a unit vector $\hat{\omega}$ giving the axis and a scalar θ giving the amount of rotation about it, combined into the single vector $\hat{\omega}\theta$ [26]. It should be kept in mind that any three-parameter description of orientation is affected by configurations at which it becomes ill-defined; this is a property of minimal representations in general, and not a shortcoming of this particular choice.

Whatever representation is chosen, the direct kinematics problem always admits one and only one solution: given the joint variables, the pose of the end-effector follows uniquely from (2.5). The inverse problem, examined next, does not share this property.

2.2.2 Differential Kinematics

Differential Kinematics (DK) describes the relationship between the joint velocities $\dot{q}$ and the velocity of the end-effector.

The velocity of the end-effector is described by two three-dimensional vectors: the linear velocity $\dot{p}_e \in \mathbb{R}^3$ of the origin of its frame and its angular velocity $\omega_e \in \mathbb{R}^3$. These are collected into the twist

$$v = \begin{bmatrix} \dot{p}_e \\ \omega_e \end{bmatrix} \in \mathbb{R}^6. \tag{2.7}$$

Each of the two contributions depends linearly on the joint velocity vector $\dot{q}$, through distinct matrices:

$$\dot{p}_e = J_P(q)\,\dot{q}, \qquad \omega_e = J_O(q)\,\dot{q}, \tag{2.8}$$

with $J_P(q), J_O(q) \in \mathbb{R}^{3\times n}$. Stacking the two blocks yields the geometric Jacobian of the chain and the compact relation

$$J(q) = \begin{bmatrix} J_P(q) \\ J_O(q) \end{bmatrix} \in \mathbb{R}^{6\times n}, \qquad v = J(q)\,\dot{q}. \tag{2.9}$$

It should be emphasized that J relates the joint velocities to the twist (2.7), and not, in general, to the time derivative of the minimal pose representation (2.6). The two are related by a further transformation depending on the orientation representation adopted, and the matrix obtained in this way is referred to as the analytical Jacobian [24].

It should also be emphasized that J depends on the current configuration: for a given joint velocity, the velocity of the end-effector varies depending on the configuration of the manipulator [24]. It is precisely this dependence that makes it necessary to recompute the Jacobian at every step in iterative algorithms.

Computation of the Jacobian

The columns of J admit a direct geometric interpretation: the i-th column represents the contribution of joint i to the end-effector velocity when all the other joints are held still. Its expression depends on the type of joint. For a revolute joint, the motion is a rotation about the joint axis, so the angular contribution is the axis itself and the linear contribution is the velocity of the end-effector origin as it rotates about that axis; for a prismatic joint, the motion is a pure translation along the axis, which produces no angular velocity. Denoting by z_{i-1} the unit vector of the axis of joint i and by p_{i-1} the position of the origin of frame $i-1$, this gives

$$J^{(i)}(q) = \begin{cases} \begin{bmatrix} z_{i-1} \times (p_e - p_{i-1}) \\ z_{i-1} \end{bmatrix} & \text{revolute joint,} \\ \begin{bmatrix} z_{i-1} \\ 0 \end{bmatrix} & \text{prismatic joint.} \end{cases} \tag{2.10}$$

Both quantities appearing in (2.10) are already available from the direct kinematics: since the columns of a homogeneous transformation matrix contain the axes of the corresponding frame and the position of its origin, z_{i-1} and p_{i-1} are obtained as the first three elements of the third and fourth column of

$$T_{i-1}^{b} = T_0^b \prod_{j=1}^{i-1} A_j^{j-1}(q_j), \tag{2.11}$$

so that the Jacobian is assembled while traversing the chain, using quantities already obtained from the forward kinematics [24].

Relation (2.9) can finally be read in the opposite direction. Given a desired end-effector velocity, one may ask which joint velocities produce it, which amounts to inverting the Jacobian. This is the step that turns differential kinematics into a practical tool for solving the inverse kinematics problem iteratively, as developed in Chapter 4.

2.2.3 Inverse Kinematics

Inverse Kinematics (IK) addresses the complementary problem: given a specific pose of the end-effector, expressed either as a matrix $T^b_{e,\star} \in SE(3)$ or as a vector $x^b_{e,\star} \in \mathbb{R}^6$, finding a joint configuration $q^\star \in \mathbb{R}^n$ that corresponds to it. The inverse kinematics problem is much more complex for three distinct reasons [24]:

- **Nonlinearity.** Equating (2.5) to the desired pose produces a system of nonlinear equations in the joint variables, for which a closed-form solution cannot be guaranteed in general. Analytical solutions are available only for particular kinematic structures, typically those in which the geometry allows position and orientation to be decoupled.
- **Multiplicity of solutions.** Different joint configurations may produce the same end-effector pose. When the number of joints exceeds the dimension of the task space, that is $n > 6$, the manipulator is kinematically redundant and the solutions are in general infinite, since internal motions exist that leave the end-effector pose unchanged.
- **Non-existence of solutions.** A pose lying outside the manipulator workspace, or incompatible with its mechanical joint limits, admits no solution at all. Existence is guaranteed only for poses belonging to the reachable workspace of the structure.

A convenient way of dealing with these difficulties is to restate the problem as the minimization of the discrepancy between the desired pose and the one produced by the current configuration. The discrepancy must first be given the form of a vector. Its positional part is simply the difference between desired and current position, whereas the orientation part requires more care, since the difference between two minimal orientation representations is not a meaningful measure of relative rotation. The standard construction forms the relative rotation between the current and the desired end-effector orientation and extracts its axis–angle representation through the matrix logarithm on $SO(3)$ [26], which yields the pose-error vector

$$e(q) = \begin{bmatrix} p^b_{e,\star} - p^b_e(q) \\ \mathrm{Log}\big(R^b_{e,\star}\, R^b_e(q)^\top\big)^\vee \end{bmatrix} \in \mathbb{R}^6, \tag{2.12}$$

where $\mathrm{Log}(\cdot)^\vee$ denotes the matrix logarithm on $SO(3)$ followed by the extraction of the associated three-dimensional vector, so that both components of e are expressed in the base frame. The orientation term of (2.12) is precisely the axis–angle representation introduced in Section 2.2.1, applied to the relative rotation $R^b_{e,\star} R^b_e(q)^\top$ that carries the current orientation onto the desired one: it is the vector $\hat{\omega}\theta$ whose direction $\hat{\omega}$ is the axis of that rotation and whose magnitude θ is the angle to be rotated about it. The orientation error is therefore expressed, like the positional one, as a three-dimensional vector that vanishes when the two orientations coincide. The inverse kinematics problem can then be written as

$$q^\star = \arg\min_{q \in \mathbb{R}^n} \tfrac{1}{2} \|e(q)\|^2. \tag{2.13}$$

Formulation (2.13) remains meaningful even when the desired pose is not exactly attainable, since it returns the configurations that approach it as closely as possible, and it accommodates additional objectives or constraints in a natural way, a feature that will be exploited extensively in the remainder of this work.

In the iterative schemes considered in the following, the joint correction is related to the pose error through the geometric Jacobian introduced in the previous section. This correspondence is exact for the translational component; for the orientation error expressed through the logarithm on $SO(3)$ it constitutes a local first-order approximation, which becomes exact as the orientation error vanishes.

The choice among the available solution strategies depends on the requirements of the application. Closed-form analytical solutions are exact and inexpensive, but they are available only for particular kinematic structures; for redundant manipulators they additionally require a criterion for selecting one solution within a family. Numerical methods are applicable to any kinematic chain and accommodate secondary criteria naturally, at the cost of an iterative solution and without returning the complete set of admissible configurations [27].

Inverting the differential relation (2.9) raises two questions that the following section addresses: what to do when J is not square, as is the case for redundant manipulators, and what to do when it loses rank.

2.3 Redundancy and damped pseudoinverse

This section addresses the two questions raised at the end of the previous section: the case in which the Jacobian is not square, so that the inversion of the differential kinematics is not uniquely defined, and the configurations at which it loses rank, where the inversion becomes numerically unreliable.

2.3.1 Kinematic redundancy

A manipulator is said to be kinematically redundant when it possesses a number of degrees of freedom greater than the number of variables required to describe the assigned task. With reference to the spaces introduced in Section 2.2, a manipulator is intrinsically redundant when the dimension m of the operational space is smaller than the dimension n of the joint space.

It is important to remark that redundancy is a property relative to the task, and not an absolute attribute of the mechanical structure: a manipulator may be redundant with respect to one task and non-redundant with respect to another. Even when $m = n$, a manipulator is said to be functionally redundant whenever only $r < m$ components of the operational space are of concern for the specific task to be executed [24].

Internal motions. When a manipulator is redundant and its Jacobian has full row rank, the null space of the Jacobian is non-trivial: there exist joint velocities that produce no motion of the end-effector at all, usually referred to as internal motions, so that the structure can reconfigure itself while the pose of the end-effector remains unchanged. Redundancy is therefore not a deficiency of the mechanism but an

additional resource, since the surplus degrees of freedom can be exploited to satisfy further requirements while the assigned task is being carried out.

2.3.2 Pseudoinverse solutions and the null space

Inverting relation (2.9) is well defined only when the Jacobian is square and non-singular. The inverse is therefore replaced by the Moore–Penrose pseudoinverse $J^\dagger$, which is defined for any matrix, square or not.

When $J \in \mathbb{R}^{m\times n}$ has more columns than rows, that is $m < n$, the system is underdetermined and infinitely many joint velocities produce the same end-effector velocity. Among these, the right pseudoinverse

$$J^\dagger = J^\top \left(JJ^\top\right)^{-1} \tag{2.14}$$

selects the one of minimum Euclidean norm, and satisfies $JJ^\dagger = I_m$, so that the desired end-effector velocity is reproduced exactly. It is important to remark that (2.14) is available only if J has full row rank, since this is the condition that guarantees the invertibility of $JJ^\top$.

When instead J has more rows than columns, that is $m > n$, the system is overdetermined and in general no exact solution exists. The left pseudoinverse

$$J^\dagger = \left(J^\top J\right)^{-1} J^\top \tag{2.15}$$

returns the joint velocity that minimizes the residual $\|J\dot{q} - v\|$, that is, the least-squares solution, and satisfies $J^\dagger J = I_n$ provided that J has full column rank [24].

The two expressions above are particular instances of a single property, which is more naturally stated in terms of the subspace reachable by the mapping. If the required end-effector velocity can be produced by some joint velocity, the pseudoinverse returns the one of smallest norm among those that produce it exactly; if it cannot, the pseudoinverse returns the joint velocity of smallest norm among those that come closest to producing it. Whether one falls in the first or in the second situation therefore depends not on the shape of the Jacobian alone, but on whether the desired velocity belongs to the subspace that the Jacobian can span in the current configuration.

This subspace, together with its counterpart in the joint space, characterizes the mapping completely. The range space $\mathcal{R}(J)$ is the set of end-effector velocities that can be generated by some joint velocity in the current configuration, while the null space $\mathcal{N}(J)$ is the set of joint velocities that produce no end-effector motion at all. Their dimensions satisfy

$$\dim \mathcal{R}(J) + \dim \mathcal{N}(J) = n, \tag{2.16}$$

independently of the rank of J. A non-trivial null space is precisely what makes the solution of the underdetermined case non-unique: the matrix $I - J^\dagger J$ projects onto $\mathcal{N}(J)$, so that adding any vector of this form to a solution leaves the end-effector velocity unchanged.

2.3.3 Kinematic singularities and damped least squares

A manipulator is said to be in a singular configuration when its Jacobian loses rank, that is, when the rank falls below the maximum value the matrix attains over the configuration space. In such a configuration a direction of end-effector motion that is locally attainable elsewhere can no longer be produced by any joint velocity, and the dimension of the range space decreases [24]. For a Jacobian that is generically full row rank, such as J_r away from singularities, expression (2.14) ceases to be applicable when that rank is lost, since $JJ^\top$ is then no longer invertible. The difficulty is not confined to exact singularities: in their neighborhood the matrix becomes ill-conditioned, and the pseudoinverse returns very large joint velocities even for small displacements of the target [27].

A standard remedy consists in relaxing the requirement that the end-effector velocity is reproduced exactly, and in seeking instead a compromise between accuracy and magnitude of the joint velocities. This is obtained by minimizing

$$\|J\dot{q} - v\|^2 + \lambda\|\dot{q}\|^2, \tag{2.17}$$

where $\lambda > 0$ is a damping factor. Imposing the first-order optimality condition and rearranging the resulting normal equations gives the damped least-squares inverse

$$J^{\#} = J^\top \left(JJ^\top + \lambda I\right)^{-1}, \tag{2.18}$$

which replaces (2.14) whenever numerical robustness is required [24, 27]. The matrix to be inverted has the dimension of the task space rather than that of the joint space, which makes (2.18) convenient for redundant chains.

The damping factor governs a trade-off. A small λ leaves the solution close to the minimum-norm one and preserves accuracy, but offers little protection when the Jacobian is ill-conditioned; a large λ bounds the joint velocities at the price of a residual error on the end-effector velocity, since $JJ^{\#} \neq I$ in general. Its value is therefore selected as a compromise.

2.4 Numerical optimization: Newton, Gauss–Newton and barrier methods

The inverse kinematics problem was formulated in (2.13) as the minimization of a scalar function of the joint variables. Since no closed-form solution is available in general, such a problem must be addressed by iterative numerical methods. The following sections describe the methods employed in this work: the Newton method for unconstrained minimization, its Gauss–Newton variant for problems exhibiting a least-squares structure, and the barrier approach for problems subject to inequality constraints [28].

2.4.1 Unconstrained minimization and Newton's method

Consider the problem of minimizing a twice continuously differentiable function Φ, expressed as

$$\min_{x \in \mathbb{R}^n} \Phi(x), \tag{2.19}$$

where $\Phi : \mathbb{R}^n \to \mathbb{R}$ is the objective function, $\nabla\Phi(x) \in \mathbb{R}^n$ denotes its gradient and $\nabla^2\Phi(x) \in \mathbb{R}^{n \times n}$ its Hessian matrix.

Iterative methods approach the solution of (2.19) by generating a sequence of estimates according to the update rule

$$x^{(k+1)} = x^{(k)} + \alpha\, d^{(k)},$$

in which $d^{(k)} \in \mathbb{R}^n$ is a search direction and $\alpha > 0$ is a scalar step length. A direction is said to be a descent direction at $x^{(k)}$ when it satisfies

$$\nabla\Phi\left(x^{(k)}\right)^\top d^{(k)} < 0. \tag{2.20}$$

Condition (2.20) is a standard requirement in line-search methods, since it ensures that a sufficiently small positive step produces a decrease of the objective [28].

Newton direction. Newton's method obtains the search direction from a local quadratic model of the objective, built from its second-order Taylor expansion around the current estimate. Minimizing that model yields the Newton direction as the solution of the linear system

$$\nabla^2\Phi\left(x^{(k)}\right) d^{(k)} = -\nabla\Phi\left(x^{(k)}\right), \tag{2.21}$$

The Newton direction is therefore obtained by solving a linear system whose coefficient matrix is the Hessian evaluated at the current estimate, and whose right-hand side is the gradient with the sign reversed.

It is important to remark that (2.21) yields a descent direction whenever the Hessian is positive definite and the gradient is nonzero. In that case the inverse Hessian is positive definite as well, and

$$\nabla\Phi^\top d^{(k)} = -\nabla\Phi^\top \left(\nabla^2\Phi\right)^{-1} \nabla\Phi < 0,$$

so that (2.20) is satisfied strictly. Positive definiteness of the Hessian is therefore the property to be verified when the descent character of a Newton-type step must be established.

The step length α is either kept fixed or selected at each iteration. A common strategy consists in starting from a trial value and reducing it repeatedly until an acceptance test is satisfied, a procedure referred to as backtracking [28].

2.4.2 Gauss–Newton methods for nonlinear least squares

A particular structure arises when the objective is a sum of squares of m scalar functions, collected in the residual vector $r(x) \in \mathbb{R}^m$. In this case the objective can be written as

$$\Phi(x) = \frac{1}{2} \sum_{j=1}^{m} r_j(x)^2 = \frac{1}{2}\, r(x)^\top r(x), \tag{2.22}$$

where $r_j(x)$ denotes the j-th component of the residual vector.

Denoting by $J_{\mathrm{res}}(x) \in \mathbb{R}^{m \times n}$ the Jacobian of the residual vector, distinguished by the subscript from the kinematic Jacobians introduced in Section 2.2, the gradient and the Hessian of (2.22) assume the form

$$\nabla \Phi = J_{\mathrm{res}}^\top r, \qquad \nabla^2 \Phi = J_{\mathrm{res}}^\top J_{\mathrm{res}} + \sum_{j=1}^{m} r_j\, \nabla^2 r_j, \tag{2.23}$$

where $\nabla^2 r_j$ is the Hessian of the j-th residual [28]. The Hessian of the objective is therefore composed of two contributions: the first depends only on first derivatives of the residuals, whereas the second requires their individual second derivatives.

Gauss–Newton direction. The Gauss–Newton method exploits this structure by discarding the second contribution in (2.23), so that the search direction is obtained from

$$J_{\mathrm{res}}^\top J_{\mathrm{res}}\, d^{(k)} = -J_{\mathrm{res}}^\top r. \tag{2.24}$$

The solvability of (2.24) depends on the rank of the residual Jacobian. If J_{res} has full column rank, $J_{\mathrm{res}}^\top J_{\mathrm{res}}$ is positive definite and the Gauss–Newton direction is uniquely defined by (2.24). If J_{res} is rank deficient, the normal equations are singular and the step must instead be interpreted through the corresponding linear least-squares problem, for instance by selecting a minimum-norm solution through a pseudoinverse.

Two advantages follow from this construction:

- the individual Hessians of the residuals need not be computed, since the approximation $J_{\mathrm{res}}^\top J_{\mathrm{res}}$ can be constructed directly from the Jacobian already required for the gradient;
- the approximation is accurate when the discarded terms are small with respect to the retained ones, which occurs in particular when the residuals are small at the solution, or when each residual is nearly a linear function of the variables, so that its second derivatives are negligible.

The same principle underlies the generalized formulation presented in Chapter 5, which retains the curvature of the objective propagated through a first-order approximation of a nonlinear mapping, while neglecting the terms that involve the second derivatives of that mapping.

2.4.3 Barrier methods for inequality constraints

The methods described above address unconstrained problems. When the variables are additionally required to satisfy a set of inequality constraints, the problem takes the form

$$\min_{x} \ \Phi(x) \quad \text{subject to} \quad c_i(x) \geq 0, \quad i \in \mathcal{I}, \tag{2.25}$$

where c_i denotes the i-th constraint function and $\mathcal{I}$ the index set of the constraints. The strictly feasible region is the set of points at which all constraints hold with strict inequality.

Barrier methods enforce the constraints implicitly, by adding to the objective a term characterized by the following properties:

- it is infinite everywhere outside the strictly feasible region;
- it is smooth inside the strictly feasible region;
- its value approaches infinity as the boundary of that region is approached.

Logarithmic barrier. The most widely used choice satisfying these requirements is the logarithmic barrier, which leads to the composite objective

$$P(x; \mu) = \Phi(x) - \mu \sum_{i \in \mathcal{I}} \log c_i(x), \tag{2.26}$$

where $\mu > 0$ is the barrier parameter, weighting the barrier term against the original objective [28]. Since the logarithmic barrier is finite only in the strictly feasible region and diverges as its boundary is approached, any minimizer of (2.26), whenever it exists, lies strictly inside the feasible region, and the constrained problem is replaced by an unconstrained one that can be addressed with the methods of the previous sections.

Chapter 3

Materials

This chapter describes the experimental equipment and the technologies employed throughout this thesis. First, the collaborative robot used to reproduce the adapted trajectories is described, together with its main characteristics and kinematic model. Then, the inertial motion capture system used to record the reference movements is introduced, with particular attention to the coordinate systems and joint-angle conventions required for the subsequent processing of the acquired data.

3.1 Franka Emika Panda

The robotic manipulator used in this study is the Franka Emika Panda, installed in the Intelligent Robotics Laboratory at the University of Pavia. The Panda is a collaborative robot designed to operate in a workspace shared with human operators.

The robot is a serial manipulator composed of seven revolute joints and has a rated payload of 3 kg [29]. Since its seven degrees of freedom exceed the six variables required to describe a general end-effector pose, the Panda is kinematically redundant with respect to the Cartesian task, as discussed in Section 2.3.1. This redundancy is handled in the trajectory computation through the pseudoinverse-based differential kinematic mapping adopted in the HJTT framework and further developed in Chapter 5.

The joint position limits and the maximum joint velocity and acceleration values are reported in Table 3.1.

Table 3.1: Joint limits and kinematic constraints of the Franka Emika Panda [1].

Name	Joint 1	Joint 2	Joint 3	Joint 4	Joint 5	Joint 6	Joint 7	Unit
$q_{\max}$	2.8973	1.7628	2.8973	−0.0698	2.8973	3.7525	2.8973	rad
$q_{\min}$	−2.8973	−1.7628	−2.8973	−3.0718	−2.8973	−0.0175	−2.8973	rad
$\dot{q}_{\max}$	2.1750	2.1750	2.1750	2.1750	2.6100	2.6100	2.6100	rad/s
$\ddot{q}_{\max}$	15	7.5	10	12.5	15	20	20	rad/s^2

The values reported in Table 3.1 refer specifically to the Franka Emika Robot (FER/Panda) and are distinguished by the manufacturer from those of the successor model. They are the limits specified for the control interface through which the

robot is commanded, and are more restrictive than the nominal figures reported in the datasheet [1, 29].

The kinematic model adopted in this work follows the Denavit–Hartenberg parameterization reported in Table 3.2, consistent with the Panda model used in the original HJTT implementation [4]. The corresponding link dimensions are also consistent with the official FER kinematic model [30].

Table 3.2: Denavit–Hartenberg parameters adopted for the nominal Panda kinematic model [4].

Joint	a (m)	d (m)	α (rad)	θ (rad)
Joint 1	0	0.333	0	θ_1
Joint 2	0	0	$-\pi/2$	θ_2
Joint 3	0	0.316	$\pi/2$	θ_3
Joint 4	0.0825	0	$\pi/2$	θ_4
Joint 5	-0.0825	0.384	$-\pi/2$	θ_5
Joint 6	0	0	$\pi/2$	θ_6
Joint 7	0.088	0	$\pi/2$	θ_7
Flange	0	0.107	0	0

Since these parameters follow the modified Denavit–Hartenberg convention introduced by Craig and adopted in the Franka kinematic description [1], rather than the standard convention introduced in Section 2.2.1, they are to be used with the corresponding transformation matrix

$$A_i^{i-1}(q_i) = \begin{bmatrix} c_{\theta_i} & -s_{\theta_i} & 0 & a_i \\ s_{\theta_i} c_{\alpha_i} & c_{\theta_i} c_{\alpha_i} & -s_{\alpha_i} & -d_i s_{\alpha_i} \\ s_{\theta_i} s_{\alpha_i} & c_{\theta_i} s_{\alpha_i} & c_{\alpha_i} & d_i c_{\alpha_i} \\ 0 & 0 & 0 & 1 \end{bmatrix}, \tag{3.1}$$

where $c_{(\cdot)}$ and $s_{(\cdot)}$ denote the cosine and sine functions. The human-arm model introduced in Section 4.1 follows instead the standard convention, so that the two kinematic chains are evaluated using separate forward-kinematics routines, each consistent with its own parameterization.

The last row of Table 3.2 does not correspond to a joint, but to the fixed transformation from the last robot link to the flange, whose nominal offset is 0.107 m [30]. The interaction with the subject does not take place at the flange itself: a custom cylindrical handle, designed and 3D-printed in the laboratory, is mounted on it and provides the grip through which the arm is guided. It is shown in Figure 3.2, and the way it is used during the exercise is described in Chapter 6.

The robot is commanded through a ROS 2-based application developed in the laboratory, which transmits the adapted trajectory to the controller as a sequence of joint position references and receives the measured joint configuration during execution. The application and the experimental workflow are described in Chapter 6.

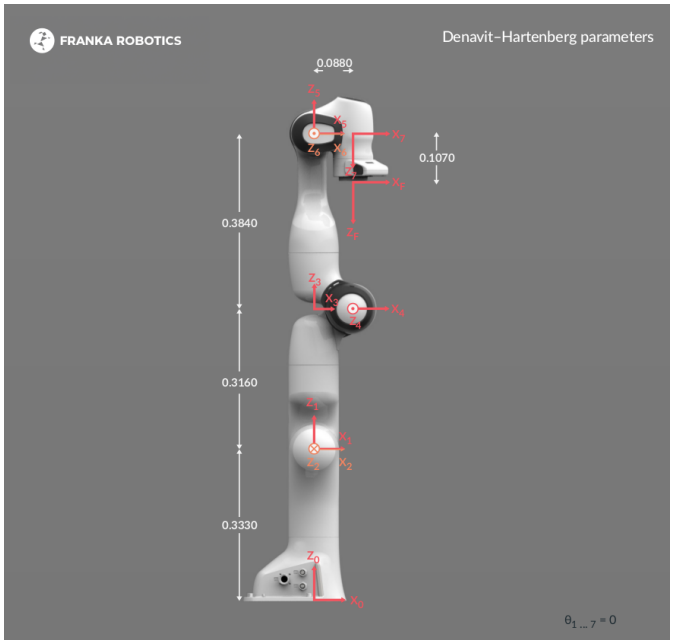


Figure 3.1: Kinematic chain of the Franka Emika Panda with the frames assigned according to the modified Denavit–Hartenberg convention, in the configuration $\theta_1 = \cdots = \theta_7 = 0$. Image from the Franka Control Interface documentation [1].



Figure 3.2: Flange of the Franka Emika Panda with the custom 3D-printed handle employed in the experimental setup. The cylindrical grip is 150 mm long and has a radius of 20 mm.

3.2 Xsens MVN Awinda

The reference human movements used in this work were acquired by means of the Xsens MVN Awinda inertial motion capture system.

Unlike optical motion capture systems, inertial systems rely on sensing units attached directly to the subject and therefore do not require external cameras or continuous line of sight. The wireless motion trackers used by the MVN Awinda system, referred to as MTw units, integrate three-dimensional accelerometers, rate gyroscopes and magnetometers, together with a barometric pressure sensor [2]. The Xsens MVN system combines the inertial measurements with a biomechanical model and sensor-fusion algorithms to estimate body-segment orientations and joint kinematics.

The configuration employed in this work is the wireless MVN Awinda system. A full-body configuration uses seventeen motion trackers mounted on the subject through dedicated straps. The trackers communicate with the Awinda Station through the Awinda wireless protocol, based on IEEE 802.15.4 technology in the 2.4 GHz frequency band. The protocol provides time synchronization of the wireless trackers to within approximately 10 μs [2].

When the Awinda Station is used as the receiver, the wireless range is approximately 20 m in indoor environments and 50 m outdoors. The corresponding range of the Awinda dongle is lower [2].



Figure 3.3: MVN Awinda system: full-body configuration worn by a subject (left) and the set of wireless MTw trackers (right). Images courtesy of Xsens Technologies B.V. [2].

An important characteristic of the system is the distinction between the internal sensor sampling, the system sample frequency and the wireless update rate. The MTw trackers internally sample angular velocity and acceleration at 1000 Hz and process these measurements on board through a Strap Down Integration (SDI) algorithm, which computes velocity and orientation increments.

According to the manufacturer, the MVN system supports a sample frequency between 60 and 240 Hz, while the wireless update rate of the Awinda configuration ranges between 20 and 60 Hz, depending on the suit configuration [2]. The use

of SDI allows accurate three-dimensional tracking to be maintained even at output rates considerably lower than the internal sampling frequency.

In the experiments performed in this work, the reference trajectories were acquired at a wireless update rate of 60 Hz.

The configuration, calibration and recording procedures are managed through the MVN Analyze/Animate software. Before acquisition, a sensor-to-segment calibration is performed to determine the alignment between the motion trackers and the corresponding anatomical body segments. The software allows the calibration to be performed in two alternative poses, shown in Figure 3.4: the N-pose, in which the subject stands upright with the arms along the body, and the T-pose, in which the arms are extended horizontally. The N-pose is the one recommended by the manufacturer and is the one adopted in all the acquisitions reported in this work. The static pose is followed by a short walk, which the manufacturer indicates as the standard calibration routine and which allows the alignment parameters to be estimated more reliably than from the static pose alone [2]. The acquired data can then be visualized, processed and exported for subsequent analysis.

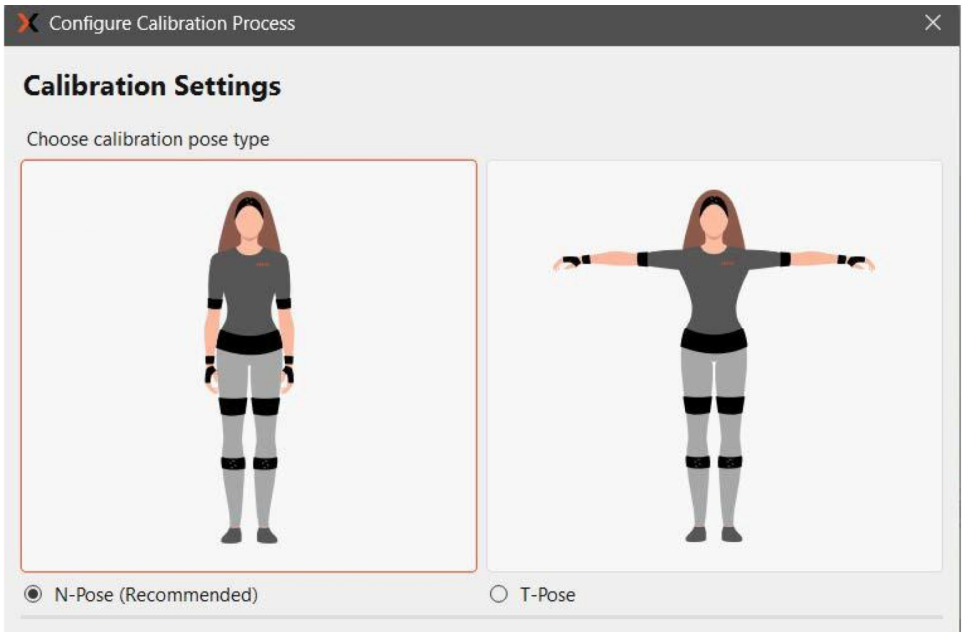


Figure 3.4: Selection of the calibration pose in MVN Analyze. The N-pose, on the left, is the pose recommended by the manufacturer and the one used in this work.

3.2.1 Coordinate systems and joint angle conventions

The MVN Fusion Engine expresses the position and orientation of each body segment with respect to an earth-fixed global reference frame G. By default, this is a right-handed Cartesian coordinate system in which the positive x-axis points towards local magnetic north, the positive y-axis points west and the positive z-axis points upwards [2]. Two distinct sets of body frames are then attached to the segments, and they serve two different purposes: the first is used to express the orientation of

each segment, the second to compute the joint angles. The distinction is relevant here because the two sets are defined through two different postures.

The frames of the first set coincide with the global frame when the subject stands in the T-pose [2]. The T-pose therefore fixes the alignment of the frames in which the segment orientations are expressed, and plays this role regardless of the pose in which the system is calibrated; in the acquisitions reported here the calibration was performed in the N-pose of Figure 3.4. After calibration, the right heel of the virtual character is placed at the origin of the scene, while the positive x-direction represents the forward direction defined during calibration.

The joint angles, which are the quantities employed in this work, are not expressed with respect to those frames. For their computation MVN introduces a second set of body frames, based on the anatomical pose and used only as an intermediate reference for this purpose [2]. Each of these frames has its origin at the proximal joint center, the x-axis pointing forward, the y-axis directed along the segment from the proximal to the distal joint and the z-axis pointing to the right of the subject; Figure 3.5 shows this orientation at the shoulder. In the neutral anatomical pose the frames of all the segments are mutually aligned: the relative rotation between consecutive segments is then the identity, and that is the configuration in which the joint angles returned by the system vanish. Their definitions follow, as closely as possible, the recommendations of the International Society of Biomechanics (ISB) [31, 2].

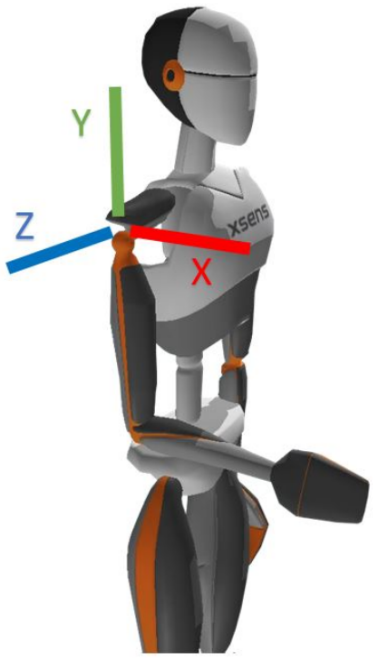


Figure 3.5: Orientation of the anatomical body frame at the right shoulder in the neutral pose. Rotations about the z-, x- and y-axis correspond respectively to flexion–extension, abduction–adduction and internal–external rotation. Adapted from [3].

Some differences from the original ISB definitions remain because MVN relies on inertial trackers and a biomechanical model rather than optical markers placed directly on anatomical landmarks. The manufacturer further notes that the shoulder is the joint for which its Euler extraction departs from the ISB recommendations.

Joint rotations are represented through Euler-angle decompositions. With the axes oriented as described above, each rotation corresponds to a functional movement of the joint,

$$Z : \text{flexion/extension}, \quad X : \text{abduction/adduction}, \quad Y : \text{internal/external rotation},$$

so that, for instance, the flexion–extension of the elbow is the rotation of the forearm with respect to the upper arm about the z-axis.

MVN provides joint-angle outputs based on both ZXY and XZY Euler sequences. The ZXY sequence is used for the standard joint-angle output, while an additional XZY representation is also exported and is commonly considered for shoulder motion [2]. The criterion indicated by the manufacturer is that the rotation with the largest expected excursion should come first in the sequence: for most joints this is the flexion–extension, whereas at the shoulder the abduction–adduction may instead be the dominant component [3].

In the present work, the ZXY representation was consistently adopted for the four joint variables employed by the Human Joint Trajectory Tracking framework: Shoulder Flexion/Extension (SFE), Shoulder Abduction/Adduction (SAA), Shoulder Internal/External Rotation (SIER) and Elbow Flexion/Extension (EFE). This choice is consistent with the criterion recalled above, since in both the movements recorded for the experiments the shoulder flexion–extension is the joint carrying the largest excursion, as reported in Section 6.1.2. This convention ensures consistency between the motion-capture data, the human-arm kinematic model and the trajectory-tracking framework developed in the following chapters. In particular, the range-of-motion bounds introduced in the following chapters are interpreted in this same joint-angle convention.

Chapter 4

Problem Statement and Modelling

This chapter formulates the problem addressed by this thesis. The kinematic model of the human arm adopted throughout the work is introduced first, since it is the element through which the algorithm adapts the recorded movement to the individual subject. The baseline Human Joint Trajectory Tracking algorithm is then recalled in the form in which it was originally proposed, with particular attention to the kinematic coupling between the arm and the manipulator and to the iterative procedure through which the robot trajectory is computed. The notion of joint range of motion is subsequently introduced and given an explicit form. The chapter closes by stating the problem that the remainder of the thesis addresses, namely the enforcement of patient-specific joint limits within the trajectory generation, together with the requirements that a solution must satisfy.

4.1 Human arm model

The algorithm considered in this work adapts the recorded movement to the subject through a parametric kinematic model of the upper limb. The model must therefore be simple enough to be identified from a small number of measurements, and detailed enough to reproduce the joint variables provided by the motion capture system.

The upper limb is modeled as an open kinematic chain composed of two rigid links, the upper arm and the forearm, connected by four revolute joints [32, 5]. Three of them are concentrated at the shoulder, which is treated as a spherical joint with three degrees of freedom, while the fourth corresponds to the elbow. The wrist is not included in the model, since the coupling with the robot occurs at the hand and the orientation of the hand is not among the quantities tracked by the algorithm.

The corresponding joint variables are collected in the vector

$$q_h = \begin{bmatrix} q_{\mathrm{SFE}} & q_{\mathrm{SAA}} & q_{\mathrm{SIER}} & q_{\mathrm{EFE}} \end{bmatrix}^\top \in Q \subset \mathbb{R}^4, \tag{4.1}$$

where SFE, SAA, SIER and EFE denote, respectively, shoulder flexion–extension, abduction–adduction, internal–external rotation and elbow flexion–extension. These are the same four anatomical quantities provided by the motion capture system in the ZXY decomposition described in Section 3.2.1. The acquired angles are converted to the sign convention adopted by the kinematic model before being

used by the algorithm, the only adjustment being a change of sign of the abduction/adduction component. The set Q collects the configurations that are admissible for the subject, and is given an explicit form in Section 4.3.

The geometry of the chain is determined by two anthropometric parameters, the length l_{UA} of the upper arm and the length l_{FA} of the forearm, measured on the subject before the session. The placement of the base frame of the chain with respect to the robot depends both on where the chair of the subject is positioned and on the height of the shoulder above the floor. The position of the chair is fixed and is not changed between sessions, so that the shoulder height remains the only quantity that varies from one subject to another. The subject-specific data required by the baseline HJTT algorithm are therefore three: the seated shoulder height and the lengths of the upper arm and of the forearm. It is through them that the same recorded movement produces different robot trajectories for different subjects.

The kinematic model is described by the Denavit–Hartenberg parameters reported in Table 4.1.

Table 4.1: Denavit–Hartenberg parameters of the human arm model, expressed in the standard convention [5, 4].

Link	θ_i [rad]	d_i [m]	a_i [m]	α_i [rad]
1 (SFE)	$q_{\mathrm{SFE}} - \pi/2$	0	0	$-\pi/2$
2 (SAA)	$q_{\mathrm{SAA}} - \pi/2$	0	0	$\pi/2$
3 (SIER)	$q_{\mathrm{SIER}} + \pi/2$	0	$-l_{UA}$	$\pi/2$
4 (EFE)	$q_{\mathrm{EFE}} - \pi/2$	l_{FA}	0	$-\pi/2$

The parameters follow the standard Denavit–Hartenberg convention introduced in Section 2.2.1; the chain is therefore evaluated with the corresponding elementary transformations. The constant offsets appearing in the θ column relate the zero of each Denavit–Hartenberg variable to the zero of the corresponding anatomical angle, so that the joint variables of the model are consistent with the motion-capture angles after the sign conversion described above.

The forward kinematics of the chain follows from (2.5), and the associated geometric Jacobian $J_h(q_h) \in \mathbb{R}^{6\times 4}$ is assembled as described in Section 2.2.2. As anticipated in Section 2.3.1, this chain possesses fewer degrees of freedom than the dimension of the operational space: J_h is therefore not square, and its inversion requires the pseudoinverse discussed in Section 2.3.2. This property has direct consequences on the coupling between the two kinematic chains, which is examined in Section 4.2.1.

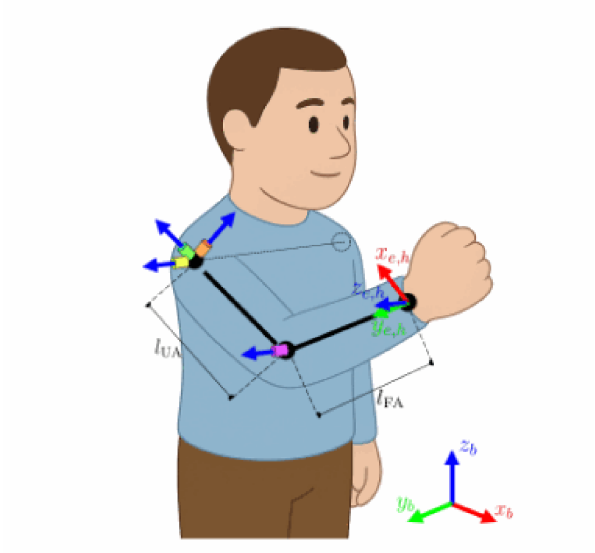


Figure 4.1: Kinematic model of the human upper limb: the rotation axes of the three shoulder joints and of the elbow, the segment lengths l_{UA} and l_{FA}, and the frame attached to the hand. Reproduced from [4].

4.2 Baseline Human Joint Trajectory Tracking Algorithm

The HJTT framework was introduced to address the limitation identified in Section 1.3.1: in end-effector-based systems the recorded reference is the Cartesian pose of the handle, and the joint configuration that the patient adopts in order to realize it depends on their own morphology [5, 4]. Rather than recording the pose, HJTT takes the joint angles of the demonstrator as the reference and computes a robot trajectory that reproduces them on the patient, exploiting the kinematic model introduced in Section 4.1.

The reference is therefore a trajectory in the human joint space,

$$q_h^\star = \begin{bmatrix} q_{\mathrm{SFE}}^\star & q_{\mathrm{SAA}}^\star & q_{\mathrm{SIER}}^\star & q_{\mathrm{EFE}}^\star \end{bmatrix}^\top,$$

acquired from the therapist by means of the motion capture system described in Section 3.2. Unlike conventional CTT, the presence of the rehabilitation robot is not required in order to define the HJTT reference: the movement is demonstrated in free space and recorded directly in the joint space of the demonstrator.

The objective of the baseline algorithm is to determine a sequence of robot configurations that produce, through the mechanical coupling with the patient, a human joint trajectory as close as possible to the one recorded from the therapist [5, 4]. Since the adaptation is based on a patient-specific kinematic model, the resulting robot trajectory does not merely reproduce the end-effector motion of the therapist, but adapts it to the morphology of the patient.

The overall structure of the framework is summarized in Figure 4.2. The movement is first demonstrated by the therapist in free space and recorded in the joint space, producing the reference $q_h^\star$. This reference and the Denavit–Hartenberg pa-

rameters of the patient are the two inputs of the algorithm, which returns the robot joint trajectory $q_r^\star$. The exercise is then executed with the patient coupled to the end-effector through the handle. Since the kinematic parameters enter the computation explicitly, the same recorded reference produces a different robot trajectory for each patient.

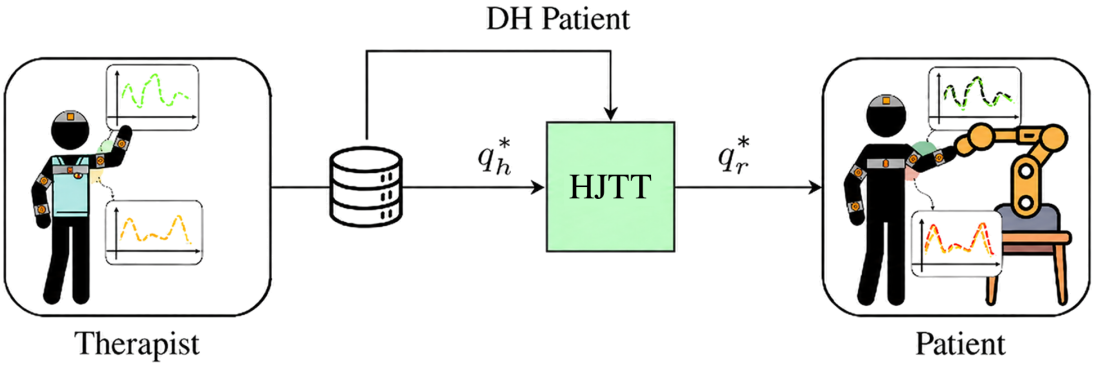


Figure 4.2: Block diagram of the baseline HJTT framework: the joint trajectory $q_h^\star$ recorded from the therapist and the Denavit–Hartenberg parameters of the patient are combined to generate the robot joint trajectory $q_r^\star$ executed during the exercise. Adapted from [5].

4.2.1 Kinematic coupling between the human arm and the robot

During a passive rehabilitation exercise, it is assumed that the patient's hand remains rigidly connected to the robot's end-effector via the handle. In end-effector based solutions, the robot and the human arm can be considered as two open kinematic chains.

Let $q_r \in \mathbb{R}^{n_r}$ denote the robot joint configuration and $q_h \in Q$ the human joint configuration. The forward kinematics of the two systems define the corresponding homogeneous transformation matrices $T^b_{e,r}(q_r)$ and $T^b_{e,h}(q_h)$, both of which are expressed with respect to a common base reference frame. The rigid connection introduced by the handle grasped by the human hand is formalized by the following assumption [5].

Assumption 4.1. There exists a constant homogeneous transformation matrix $T^r_h \in SE(3)$ such that

$$T^b_{e,r}\big(q_r(t)\big) = T^r_h \, T^b_{e,h}\big(q_h(t)\big) \tag{4.2}$$

for all $t \geq 0$, $q_r \in \mathbb{R}^{n_r}$ and $q_h \in Q$.

The matrix T^r_h represents the fixed geometric relationship imposed by the handle and it is assumed to remain constant throughout the exercise. The assumption is

reasonable in the clinical setting for which the method is intended: in patients with neurological impairment the hand is usually secured to the handle by means of an anti-spasticity glove, so that the grasp is maintained throughout the exercise and the relative placement of hand and end-effector does not change.

The differential relationship between the two chains can be obtained by differentiating the kinematic constraint (4.2) with respect to time. Denoting the robotic geometric Jacobian by $J_r(q_r) \in \mathbb{R}^{6\times n_r}$ and the human-arm geometric Jacobian by $J_h(q_h) \in \mathbb{R}^{6\times 4}$, the equality of the end-effector velocities yields

$$J_r(q_r)\,\dot{q}_r = J_h(q_h)\,\dot{q}_h. \tag{4.3}$$

The instantaneous variation of the human joint configuration induced by a robot motion can be written using the pseudoinverse of the human Jacobian as follows [5]:

$$\dot{q}_h = J_h^{\dagger}(q_h)\,J_r(q_r)\,\dot{q}_r. \tag{4.4}$$

The matrix appearing in (4.4) is the partial derivative of the human joint configuration with respect to the robot one,

$$\frac{\partial q_h}{\partial q_r} = J_h^{\dagger}(q_h)\,J_r(q_r) \in \mathbb{R}^{4\times n_r}, \tag{4.5}$$

and provides the local differential mapping from variations in the robot joint coordinates to the corresponding variations in the human joint coordinates. This constitutes the fundamental link exploited by HJTT to predict how a modification of the robot configuration affects the posture of the patient's arm.

It is worth remarking why a pseudoinverse is required in (4.4). The human Jacobian maps four joint velocities into a six-dimensional task space, so that the system is overdetermined and admits no exact solution in general: the human arm cannot reproduce an arbitrary end-effector motion. Consistently with the discussion of Section 2.3.2, the left pseudoinverse therefore returns the least-squares solution, and the coupling (4.3) is satisfied exactly only when the velocity imposed by the robot belongs to the range space of the human Jacobian.

4.2.2 Human joint tracking problem

The human joint tracking error is defined as the difference between the therapist's joint space configuration $q_h^{\star}$ and the patient's joint space configuration:

$$e_h = q_h^{\star} - q_h, \qquad e_h^{(k)} = q_h^{\star} - q_h^{(k)} \ \text{ at iteration } k. \tag{4.6}$$

Given a desired human joint configuration $q_h^{\star}$, the HJTT problem is to determine a robot configuration $q_r^{\star}$ such that the induced human configuration is as close as possible to the desired one:

$$q_r^{\star} = \arg\min_{q_r\in\mathbb{R}^{n_r}} \big\|q_h^{\star} - q_h(q_r)\big\|, \tag{4.7}$$

where $q_h(q_r)$ denotes the human joint configuration induced by the robot through coupled kinematics. The quantity that is minimized is the human joint tracking error rather than the pose error of the robot end-effector.

As a closed-form expression of $q_h(q_r)$ is not available, the mapping is locally approximated using differential kinematics. Around the current configurations $q_r^{(k)}$ and $q_h^{(k)}$, a first-order Taylor expansion yields

$$q_h(q_r) \simeq q_h^{(k)} + \frac{\partial q_h}{\partial q_r}\left(q_r - q_r^{(k)}\right), \qquad \frac{\partial q_h}{\partial q_r} = J_h^{\dagger}\left(q_h^{(k)}\right) J_r\left(q_r^{(k)}\right). \tag{4.8}$$

The Jacobian of the mapping follows directly from (4.4). Substituting (4.8) into the error definition (4.6), the linearized human joint error becomes

$$e_h(q_r) \simeq e_h^{(k)} - J_h^{\dagger}\left(q_h^{(k)}\right) J_r\left(q_r^{(k)}\right)\left(q_r - q_r^{(k)}\right), \tag{4.9}$$

which characterizes the local sensitivity of the human joint error to a displacement of the robot configuration.

4.2.3 Iterative robot update

The robot configuration is modified iteratively according to the correction structure adopted in the original HJTT formulation [5, 4]. The baseline HJTT correction is constructed in two stages. The human joint-space error is first mapped into an end-effector differential displacement through the human Jacobian, and this displacement is then mapped into the robot joint space through the pseudoinverse of the robot Jacobian:

$$\Delta q_r^{(k)} = J_r^{\dagger}\left(q_r^{(k)}\right) J_h\left(q_h^{(k)}\right) e_h^{(k)}. \tag{4.10}$$

The linearization (4.8)–(4.9) describes the local sensitivity of the human configuration to a robot displacement and motivates this construction; it does not, however, yield (4.10) by direct inversion, since the pseudoinverse of the composite mapping $J_h^{\dagger} J_r$ does not coincide in general with $J_r^{\dagger} J_h$.

The iterative update rule is therefore

$$\boxed{q_r^{(k+1)} = q_r^{(k)} + \alpha\, J_r^{\dagger}\left(q_r^{(k)}\right) J_h\left(q_h^{(k)}\right)\left(q_h^{\star} - q_h^{(k)}\right)} \tag{4.11}$$

where $0 < \alpha \leq 1$ is a scalar step-size parameter that controls the amplitude of the iterative correction and contributes to the stability of the convergence. In the reference implementation, the full correction is applied, corresponding to $\alpha = 1$.

Compactly, the correction passes through the task space,

$$e_h \xrightarrow{J_h} \Delta x \xrightarrow{J_r^{\dagger}} \Delta q_r,$$

so that the reference remains defined in the human joint space, while the actual correction is realized through the redundant robot kinematics.

Since the manipulator may approach singular or ill-conditioned configurations, the implementation does not rely on an undamped Moore–Penrose pseudoinverse. Instead, a Damped Least-Squares (DLS) inverse is adopted:

$$J^{\#} = J^{\top}\left(JJ^{\top} + \lambda I\right)^{-1}, \tag{4.12}$$

with $\lambda > 0$ the damping factor, whose purpose is to improve the numerical robustness of the inverse differential kinematics in the neighborhood of singular configurations [24, 27]. This damped form is applied to both kinematic chains, that is $J_r^{\#} = J_r^{\top}(J_r J_r^{\top} + \lambda I)^{-1}$ for the robot and $J_h^{\#} = J_h^{\top}(J_h J_h^{\top} + \lambda I)^{-1}$ for the human arm model. Accordingly, $J_r^{\dagger}$ and $J_h^{\dagger}$ in (4.10)–(4.11) are replaced by $J_r^{\#}$ and $J_h^{\#}$ in the implemented formulation.

The damping is what makes the right-regularized form applicable to the arm model as well. Its Jacobian $J_h \in \mathbb{R}^{6\times 4}$ has rank at most four, so the product $J_h J_h^{\top}$ is a 6×6 matrix of rank at most four and is therefore singular in every configuration: without the damping term, (4.12) would not be defined for this chain. This should not be interpreted as a condition for the existence of the Moore–Penrose pseudoinverse itself, which for the arm model is given by the left form wherever J_h has full column rank. Singular configurations of the arm model are those in which that rank drops below four, and are not the generic case.

4.2.4 Update of the human configuration

Once the robot configuration has been updated, the corresponding variation of the human posture must also be estimated. This is obtained by integrating the local differential relationship (4.4) over the robot increment:

$$q_h^{(k+1)} = q_h^{(k)} + J_h^{\#}\left(q_h^{(k)}\right) J_r\left(q_r^{(k)}\right)\left(q_r^{(k+1)} - q_r^{(k)}\right). \tag{4.13}$$

The updated human configuration is therefore not obtained by solving an independent inverse-kinematics problem at every inner iteration, but is propagated from the current estimate through the differential robot-to-human mapping. The human and robot configurations consequently evolve together during the iterative procedure.

At each iteration, starting from $q_r^{(k)}$ and $q_h^{(k)}$, the forward kinematics and the geometric Jacobians of both chains are evaluated, the current error $e_h^{(k)}$ is computed and used to determine the new robot configuration through (4.11), and the resulting robot displacement is mapped back to the human joint space through (4.13). The procedure is repeated until the convergence condition

$$\left\| q_h^{\star} - q_h^{(k)} \right\| < \varepsilon_{\text{conv}} \tag{4.14}$$

is satisfied, where $\varepsilon_{\text{conv}} > 0$ is a prescribed tolerance. A maximum number of iterations $k_{\max}$ is also specified in order to avoid an indefinite optimization loop when convergence cannot be reached. The reference implementation adopts $\varepsilon_{\text{conv}} = 10^{-3}$ and $k_{\max} = 10^4$.

The procedure is summarized in Algorithm 4.1, where H_{DH} and R_{DH} denote the sets of Denavit–Hartenberg parameters of the human arm model and of the manipulator, reported in Table 4.1 and Table 3.2 respectively. Its outputs are the robot configuration $q_r^{\star}$ that solves (4.7) and the corresponding human configuration $q_h^{\star\star}$ predicted by the kinematic model, which is distinguished from the reference $q_h^{\star}$ because the two coincide only when the iteration converges.

Algorithm 4.1 HJTT optimization procedure for a given desired configuration $q_h^\star$

Require: $q_h^\star \in Q$, $q_h^{(0)} \in Q$, $q_r^{(0)} \in \mathbb{R}^{n_r}$, H_{DH}, R_{DH}, $\alpha \in (0,1]$, $\lambda > 0$, $\varepsilon_{\text{conv}} > 0$, $k_{\max}$

Ensure: $q_r^\star$, $q_h^{\star\star}$, δ_{conv}

$k \leftarrow 0$, $\delta_{\text{conv}} \leftarrow$ FALSE

while $\delta_{\text{conv}} =$ FALSE **and** $k < k_{\max}$ **do**

  compute the human forward kinematics from H_{DH} and $q_h^{(k)}$ and assemble $J_h(q_h^{(k)})$

  compute the robot forward kinematics from R_{DH} and $q_r^{(k)}$ and assemble $J_r(q_r^{(k)})$

  $J_h^\# \leftarrow J_h^\top (J_h J_h^\top + \lambda I)^{-1}$, $J_r^\# \leftarrow J_r^\top (J_r J_r^\top + \lambda I)^{-1}$ ▷ eq. (4.12)

  $e_h^{(k)} \leftarrow q_h^\star - q_h^{(k)}$ ▷ eq. (4.6)

  $q_r^{(k+1)} \leftarrow q_r^{(k)} + \alpha\, J_r^\# J_h\, e_h^{(k)}$ ▷ eq. (4.11)

  $q_h^{(k+1)} \leftarrow q_h^{(k)} + J_h^\# J_r (q_r^{(k+1)} - q_r^{(k)})$ ▷ eq. (4.13)

  $\delta_{\text{conv}} \leftarrow \big(\|e_h^{(k)}\| < \varepsilon_{\text{conv}}\big)$ ▷ eq. (4.14)

  $k \leftarrow k + 1$

end while

$q_r^\star \leftarrow q_r^{(k)}$, $q_h^{\star\star} \leftarrow q_h^{(k)}$

4.2.5 Application to a complete trajectory

The formulation described above refers to a single desired human joint configuration, whereas a rehabilitation exercise consists of a complete sequence $\{q_h^\star(t_i)\}_{i=1}^N$. The HJTT procedure is therefore applied sequentially to all samples of the recorded trajectory. For each time instant t_i, the algorithm determines a robot configuration $q_r^\star(t_i)$ which minimizes the human joint-space discrepancy. The final result is a patient-adapted robot trajectory $\{q_r^\star(t_i)\}_{i=1}^N$.

Initialization plays an important role in obtaining a continuous sequence of solutions. As no previous robot configuration is available for the first sample, the robot is initialized by solving a standard inverse kinematics problem [24], in which the target position is the one obtained from the forward kinematics of the human arm model in the first desired configuration, and the target orientation is a prescribed end-effector orientation:

$$p_{e,r}^b\big(q_r^\star(t_0)\big) = p_{e,h}^b\big(q_h^\star(t_0)\big), \qquad R_{e,r}^b\big(q_r^\star(t_0)\big) = \bar{R}, \tag{4.15}$$

where $\bar{R}$ denotes the fixed orientation adopted for the initial pose. This provides a consistent initial condition for the two kinematic chains. If the inverse kinematics solver does not converge, a warning is issued and the algorithm proceeds using the best available approximation.

For all subsequent samples, the previously computed robot configuration is used as the robot initial guess, while the human configuration is initialized from the previous sample of the desired trajectory,

$$q_r^{(0)}(t_i) = q_r^\star(t_{i-1}), \qquad q_h^{(0)}(t_i) = q_h^\star(t_{i-1}), \qquad i > 0.$$

This reduces the distance between consecutive optimization problems and promotes continuity along the computed trajectory. The iterative procedure is executed for each sample until either the convergence condition (4.14) is satisfied or $k_{\max}$ is reached.

The baseline HJTT can therefore be summarized as a patient-specific inverse-kinematics procedure whose reference is defined in the human joint space: the therapist provides the desired joint trajectory, the patient-specific arm model determines the coupling between human and robot kinematics, and an iterative differential procedure computes the robot trajectory required to reproduce the desired joint behavior. The trajectory is generated offline and subsequently executed to guide the patient through the movement. This kinematic framework forms the basis for the ROM-constrained formulation introduced in the following sections.

4.3 Human Joint Range of Motion and Patient-Specific Limits

The range of motion of a joint is the angular excursion that the joint can perform, and it constitutes one of the fundamental parameters used in clinical practice to characterize motor function. Even among healthy individuals, ROM is not a fixed quantity, as it varies systematically with factors such as age and sex and may differ from commonly used normative values [33]. In patients undergoing rehabilitation, the admissible excursion is generally further reduced, for the reasons recalled in Section 1.4 [20], and may evolve over the course of treatment. The ROM is therefore a patient-specific quantity, which reflects the clinical condition of the individual rather than a general anatomical property.

Denoting by $\mathcal{C} \subseteq \{1, \dots, 4\}$ the set of joints for which a joint limitation is specified, and by $q_{h,\min,i}$ and $q_{h,\max,i}$ the corresponding lower and upper bounds, the admissible set introduced in Section 4.1 can be given the explicit form

$$Q = \left\{ q_h \in \mathbb{R}^4 \ : \ q_{h,\min,i} \leq q_{h,i} \leq q_{h,\max,i}, \ i \in \mathcal{C} \right\}. \tag{4.16}$$

The set (4.16) is a box-shaped parameterization of the admissible configurations, in which each joint is bounded independently of the others. This representation is the one most commonly adopted in clinical practice, since the ROM of each degree of freedom is measured and prescribed separately.

It should be noted that (4.16) is an approximation of the actual admissible region. The physiological ROM of a joint depends in general on the configuration of the remaining joints, owing to inter-joint dependencies, to the coupling between degrees of freedom belonging to the same articulation and to self-collision constraints; consequently, some configurations contained in the box may in fact be unfeasible for the subject. Data-driven approaches have been proposed in the literature to learn more realistic, coupled boundaries of the admissible joint space [20], and joint limits are treated as explicit constraints in optimization-based models of human motion [32]. The box representation (4.16) is nevertheless retained in this work as a first approximation, consistent with clinical practice and sufficient to formulate and validate the proposed constrained algorithm.

The relevance of ROM limits is particularly critical in the passive modality considered in this work. As observed in Section 1.4, the patient does not voluntarily contribute to the movement and cannot spontaneously oppose a motion that approaches the boundary of the admissible range: the responsibility of remaining within the physiological limits therefore rests entirely on the trajectory imposed by the robotic system. Exceeding those limits may cause pain, tissue stress or injury, and would compromise the therapeutic value of the exercise.

4.4 Problem Statement: Enforcing Patient-Specific ROM Constraints

The baseline formulation acknowledges the existence of an admissible set for the human configuration, but treats membership in that set as an assumption on the data rather than as an explicit constraint on the iterative solution: the minimization (4.7) is performed over $q_r \in \mathbb{R}^{n_r}$ without restrictions, and no mechanism in the iteration of Algorithm 4.1 prevents the estimated configuration $q_h^{(k)}$ from leaving the admissible set.

In the original setting, the reference trajectory is acquired from the therapist and assumed to be physiologically admissible. This assumption is implicitly satisfied in the scenario for which the baseline algorithm was validated, in which both the therapist and the subject acting as a patient are healthy [4]. Once a patient-specific ROM is considered, however, the therapist reference need not belong to the admissible set Q_{patient} of the patient. Consider the case in which the joint ROM of the patient is reduced with respect to that of the person who demonstrated the movement,

$$Q_{\text{patient}} \subset Q_{\text{therapist}}. \tag{4.17}$$

If the demonstrated trajectory enters the region $Q_{\text{therapist}} \setminus Q_{\text{patient}}$, there are time instants at which $q_h^\star(t) \notin Q_{\text{patient}}$. At those instants the baseline algorithm drives the patient's joints towards a configuration that lies outside their physiological range, because it minimizes the tracking error without any notion of admissibility. The more the therapist's movement exceeds the capabilities of the patient, the larger the resulting violation.

The problem addressed in this thesis can therefore be stated as follows. Given a reference trajectory $\{q_h^\star(t_i)\}_{i=1}^N$ recorded from the therapist and the admissible set Q of the patient, determine a robot trajectory $\{q_r(t_i)\}_{i=1}^N$ such that the induced human configuration reproduces the reference as closely as possible while remaining inside the admissible set at all times:

$$\min_{q_r} \ \left\| q_h^\star - q_h(q_r) \right\| \quad \text{subject to} \quad q_h(q_r) \in Q. \tag{4.18}$$

Problem (4.18) is the constrained counterpart of (4.7). Its solution must satisfy a number of requirements that follow from the structure of the baseline algorithm and from the nature of the application.

First, the constraint must be enforced within the iterative procedure and not applied as a post-processing operation on the computed trajectory. Saturating the

human configuration a posteriori would break the kinematic consistency established by (4.4) between the two chains: the resulting pair (q_r, q_h) would no longer satisfy the coupling relation, and the robot trajectory would no longer correspond to the modified human configuration.

Second, the redundancy of the manipulator must be resolved consistently with the coupling condition of Assumption 4.1, so that the robot end-effector continues to follow the hand of the patient throughout the exercise.

Third, the algorithm must degrade gracefully when the reference is not attainable. Since $q_h^\star(t) \notin Q$ occurs whenever the demonstrated movement exceeds the capabilities of the patient, the procedure cannot be expected to reduce the tracking error to zero at those instants: it must instead seek a feasible configuration that remains inside Q while limiting the deviation from the reference, rather than diverging or failing.

Finally, the constrained formulation must reduce to the baseline HJTT when the ROM term is deactivated. For feasible references sufficiently far from the ROM boundaries, the influence of the barrier should remain limited.

The formulation proposed in the following chapter addresses Problem (4.18) by augmenting the cost function with a logarithmic barrier associated with the ROM limits, and by solving the resulting optimization problem with a Generalized Gauss–Newton scheme whose correction is realized on the robot through the same differential coupling adopted by the baseline algorithm.

Chapter 5

ROM-Constrained HJTT: Proposed Method

This chapter develops the formulation announced at the end of Chapter 4 and derives the resulting iterative procedure.

Two consequences of the barrier guide the entire derivation. Firstly, the minimizer of the cost function no longer corresponds to a zero of the tracking error, so that the correction (4.10), which is driven by e_h alone, no longer points towards the solution. Secondly, the barrier is defined only within Q, so admissibility must be preserved at every iteration rather than being checked a posteriori.

5.1 Cost Function with Logarithmic Barrier

Let $\mathcal{C}$ be the set of constrained joints and let r_i, a_i and b_i be, respectively, the amplitude of the i-th ROM and the distances of the current configuration from its lower and upper bound, as introduced in (4.16):

$$r_i = q_{h,\max,i} - q_{h,\min,i}, \qquad a_i = q_{h,i} - q_{h,\min,i}, \qquad b_i = q_{h,\max,i} - q_{h,i}. \tag{5.1}$$

Introducing the normalized distances $\bar{a}_i = a_i/r_i$ and $\bar{b}_i = b_i/r_i$, the constrained cost function is defined as

$$\Phi(q_h) = \frac{1}{2} e_h^\top e_h + \gamma \sum_{i \in \mathcal{C}} r_i^2 \left[-\log \bar{a}_i - \log \bar{b}_i \right], \qquad \gamma > 0, \tag{5.2}$$

where e_h is the tracking error defined in (4.6) and the logarithmic term is the barrier function introduced in Section 2.4.3 [28]. The first term of (5.2) preserves the objective of the baseline algorithm, while the logarithmic terms diverge as the configuration approaches either bound of the admissible set. The domain of Φ is therefore the interior of Q, characterized by $a_i > 0$ and $b_i > 0$.

The amplitude r_i appears in equation (5.2) in two distinct roles. Inside the logarithm it renders the argument dimensionless and, being constant, contributes only an additive term that vanishes upon differentiation. The factor r_i^2 outside the

logarithm, conversely, survives differentiation and normalizes the curvature contribution of the barrier across joints with different ROM widths; its role is examined in Section 5.5.

Since the barrier contributes over the entire admissible domain, the minimizer of (5.2) differs from $q_h^\star$ even when the reference is attainable: the solution balances tracking accuracy with the distance from the boundaries. The magnitude of this bias is determined by γ and is discussed in Section 5.6.

5.2 Gradient and Hessian in the Human Joint Space

Differentiating the quadratic term of (5.2) gives $-e_h$, as in the baseline formulation. For the i-th barrier contribution,

$$\frac{\partial}{\partial q_{h,i}}\left[-\gamma r_i^2 \log \bar{a}_i\right] = -\frac{\gamma r_i^2}{a_i}, \qquad \frac{\partial}{\partial q_{h,i}}\left[-\gamma r_i^2 \log \bar{b}_i\right] = +\frac{\gamma r_i^2}{b_i}. \tag{5.3}$$

Collecting both contributions, the gradient admits the compact form

$$\nabla_{q_h}\Phi = -\,\tilde{e}_h, \tag{5.4}$$

where the modified error $\tilde{e}_h$ is defined component-wise as

$$\tilde{e}_{h,i} = e_{h,i} + \gamma\, r_i^2 \left(\frac{1}{a_i} - \frac{1}{b_i}\right),\ i \in \mathcal{C}, \qquad \tilde{e}_{h,i} = e_{h,i},\ i \notin \mathcal{C}. \tag{5.5}$$

Expression (5.5) shows that the barrier modifies the tracking error of the baseline algorithm by adding a repulsive component that pushes away from the bounds. It also shows that the unconstrained joints are unaffected.

Differentiating (5.4) once more, and observing that each barrier term depends on a single joint variable, the Hessian in the human joint space is diagonal:

$$H_h = \nabla^2_{q_h}\Phi = I + \gamma D, \qquad D = \mathrm{diag}(d_1, \ldots, d_4), \qquad d_i = r_i^2\left(\frac{1}{a_i^2} + \frac{1}{b_i^2}\right), \tag{5.6}$$

with $d_i = 0$ for $i \notin \mathcal{C}$. The contribution of the barrier to the curvature is positive, consistently with the convexity of the logarithmic terms; since $a_i, b_i > 0$ in the interior of Q, it follows that $D \succeq 0$ and hence

$$H_h = I + \gamma D \succ 0. \tag{5.7}$$

The cost function is therefore strictly convex with respect to q_h over the domain of the barrier, and H_h, being diagonal, is inverted at negligible cost and without any regularization.

5.3 Local HJTT Mapping and Generalized Gauss–Newton Approximation

The cost function is defined in the human joint space, whereas the algorithm acts on the robot configuration. The two are related by the differential coupling (4.3)–(4.4). Evaluating that relation at the current iterate defines the local kinematic sensitivity

$$J^{(k)} = J_h^{\dagger}\Big(q_h^{(k)}\Big)\, J_r\big(q_r^{(k)}\big) \in \mathbb{R}^{4\times n_r}, \qquad dq_h = J^{(k)} dq_r. \tag{5.8}$$

Holding $J^{(k)}$ fixed over the step, the predicted human configuration is affine in the robot increment, and composing it with (5.2) yields the local model

$$m_k(\Delta q_r) = \Phi\left(q_h^{(k)} + J^{(k)}\Delta q_r\right). \tag{5.9}$$

The model (5.9) is determined by the current state and by $J^{(k)}$ alone, and therefore inherits the local, path-dependent character of the baseline algorithm without requiring a closed-form relation $q_h(q_r)$.

Applying the chain rule to (5.9) and evaluating at $\Delta q_r = 0$, the gradient of the local model is

$$g^{(k)} := \nabla_{\Delta q_r} m_k(0) = -\left(J^{(k)}\right)^{\top} \tilde{e}_h^{(k)}, \tag{5.10}$$

which reproduces the $J^{\top}e$ structure characteristic of Gauss–Newton methods [28], with the barrier embedded in the modified error. Retaining the dependence of the sensitivity on the configuration, the variation of the robot-space gradient would read

$$dg = \left(J^{\top} H_h J + H_{\text{kin}}\right) dq_r, \tag{5.11}$$

where H_{kin} collects the terms generated by the variation of J, namely the derivatives of J_r with respect to q_r and those of $J_h^{\dagger}$ induced by the corresponding variation of q_h. These terms are equivalent to second derivatives of the kinematic relations and are computationally expensive to evaluate.

The Generalized Gauss–Newton approximation consists in neglecting H_{kin} while retaining the pullback of the human-space curvature [28]:

$$H_{\text{GGN}}^{(k)} = \left(J^{(k)}\right)^{\top} H_h^{(k)} J^{(k)} = \left(J^{(k)}\right)^{\top} J^{(k)} + \gamma \left(J^{(k)}\right)^{\top} D^{(k)} J^{(k)}, \tag{5.12}$$

in which the first term derives from the tracking objective and the second from the barrier. The term discarded by the approximation is thus the kinematic one; the curvature of the barrier, being diagonal and inexpensive, is entirely retained within H_h.

Matrix H_{GGN} in (5.12) is symmetric and positive semidefinite, but structurally rank deficient whenever the manipulator is redundant with respect to the human model:

$$\operatorname{rank}\Big(H_{\text{GGN}}^{(k)}\Big) \le \operatorname{rank}\big(J^{(k)}\big) \le 4 < n_r. \tag{5.13}$$

The positive definiteness established in (5.7) does not remove this deficiency, which originates from the dimensions of $J^{(k)}$ rather than from the conditioning of H_h. For the manipulator considered in this work, $n_r = 7$, and the null space of $H^{(k)}_{\mathrm{GGN}}$ has dimension at least three, corresponding to robot joint increments that are not observable through the local human-joint mapping $J^{(k)}$.

5.4 Gauss–Newton Iterative Update and Robotic Realization

The Gauss–Newton increment follows from setting to zero the first-order expansion of the gradient of the local model about $\Delta q_r = 0$, that is $g^{(k)} + H^{(k)}_{\mathrm{GGN}} \Delta q_r = 0$, which by (5.10) gives the normal equations

$$\left(J^{(k)}\right)^{\top} \left(I + \gamma D^{(k)}\right) J^{(k)} \Delta q_r = \left(J^{(k)}\right)^{\top} \tilde{e}^{(k)}_h. \tag{5.14}$$

By virtue of (5.13), the system (5.14) does not define a unique increment and requires a redundancy resolution criterion. Adding a robot-space regularization term would restore invertibility [28], but it would introduce an additional norm criterion that is not part of the original human-space objective; in the vanishing-regularization limit, such a formulation selects the minimum-norm robot-space solution. That criterion is inappropriate here: the cost function does not depend on the null space of the manipulator, and the minimum-norm increment may displace the end-effector along task directions outside $\mathcal{R}(J_h)$, that is, directions that the human arm cannot reproduce. Since the end-effector is rigidly coupled to the hand of the patient by Assumption 4.1, redundancy must instead be resolved consistently with that coupling.

The adopted realization separates the two problems. The correction prescribed by the cost function is computed in the human joint space, where the Hessian is diagonal and positive definite:

$$\Delta q^{(k)}_h = \left(I + \gamma D^{(k)}\right)^{-1} \tilde{e}^{(k)}_h, \tag{5.15}$$

and is then mapped onto the redundant robot configuration through the damped least-squares inverse (4.12):

$$\Delta q^{(k)}_r = J^{\#}_r\left(q^{(k)}_r\right) J_h\left(q^{(k)}_h\right) \Delta q^{(k)}_h. \tag{5.16}$$

In the undamped full-rank limit, the resulting increment is a solution of (5.14). Indeed, if $J_r J_r^{\dagger} = I_6$ and $J_h^{\dagger} J_h = I_4$, then $J^{(k)} \Delta q^{(k)}_r = \Delta q^{(k)}_h$, and substituting into the left-hand side of (5.14) gives

$$\left(J^{(k)}\right)^{\top} H^{(k)}_h J^{(k)} \Delta q^{(k)}_r = \left(J^{(k)}\right)^{\top} H^{(k)}_h \left(H^{(k)}_h\right)^{-1} \tilde{e}^{(k)}_h = \left(J^{(k)}\right)^{\top} \tilde{e}^{(k)}_h. \tag{5.17}$$

In the undamped full-rank limit, the proposed increment is therefore one of the infinitely many solutions of the singular system, selected by the requirement that the kinematic coupling $J_r \Delta q_r = J_h \Delta q_h$ be preserved rather than by a norm criterion.

Consequently, in this limit, the resulting end-effector displacement belongs to $\mathcal{R}(J_h)$. With the damped inverses actually used in the implementation this equality is no longer exact: (5.16) provides a regularized kinematic realization of the human-space correction, preserving the same coupling structure as the baseline HJTT.

No regularization of the Gauss–Newton system is thus introduced: the rank-deficient matrix (5.12) is never assembled or inverted, and the only damping employed is the factor λ of (4.12), which acts on the kinematic inversion and provides robustness in the neighborhood of the manipulator singularities, as in the baseline algorithm.

The human-space increment (5.15) is a descent direction for the cost function. Since $\left(I + \gamma D^{(k)}\right)^{-1} \succ 0$, for $\tilde{e}_h^{(k)} \neq 0$ relations (5.4) and (5.15) give

$$\left(\nabla_{q_h}\Phi\right)^{\top} \Delta q_h^{(k)} = -\left(\tilde{e}_h^{(k)}\right)^{\top} \left(I + \gamma D^{(k)}\right)^{-1} \tilde{e}_h^{(k)} < 0. \tag{5.18}$$

The configurations are finally updated as

$$q_r^{(k+1)} = q_r^{(k)} + \alpha_k\, \Delta q_r^{(k)}, \qquad q_h^{(k+1)} = q_h^{(k)} + J_h^{\#}\left(q_h^{(k)}\right) J_r\left(q_r^{(k)}\right) \left(q_r^{(k+1)} - q_r^{(k)}\right), \tag{5.19}$$

where the human configuration is propagated exactly as in (4.13), and the step size α_k is selected as described in Section 5.6.

The complete iteration is summarized in Figure 5.1. The blocks drawn in blue are those introduced by the constrained formulation: the projection of the initial configuration into the interior of the admissible set, the evaluation of the ROM distances and of the barrier curvature, the modified error and the stopping test defined on it, the human joint-space correction (5.15), the saturation of the robot increment and the feasibility test with the associated step reduction. The remaining blocks are the steps already present in the baseline algorithm of Chapter 4, which the constrained formulation leaves unchanged.

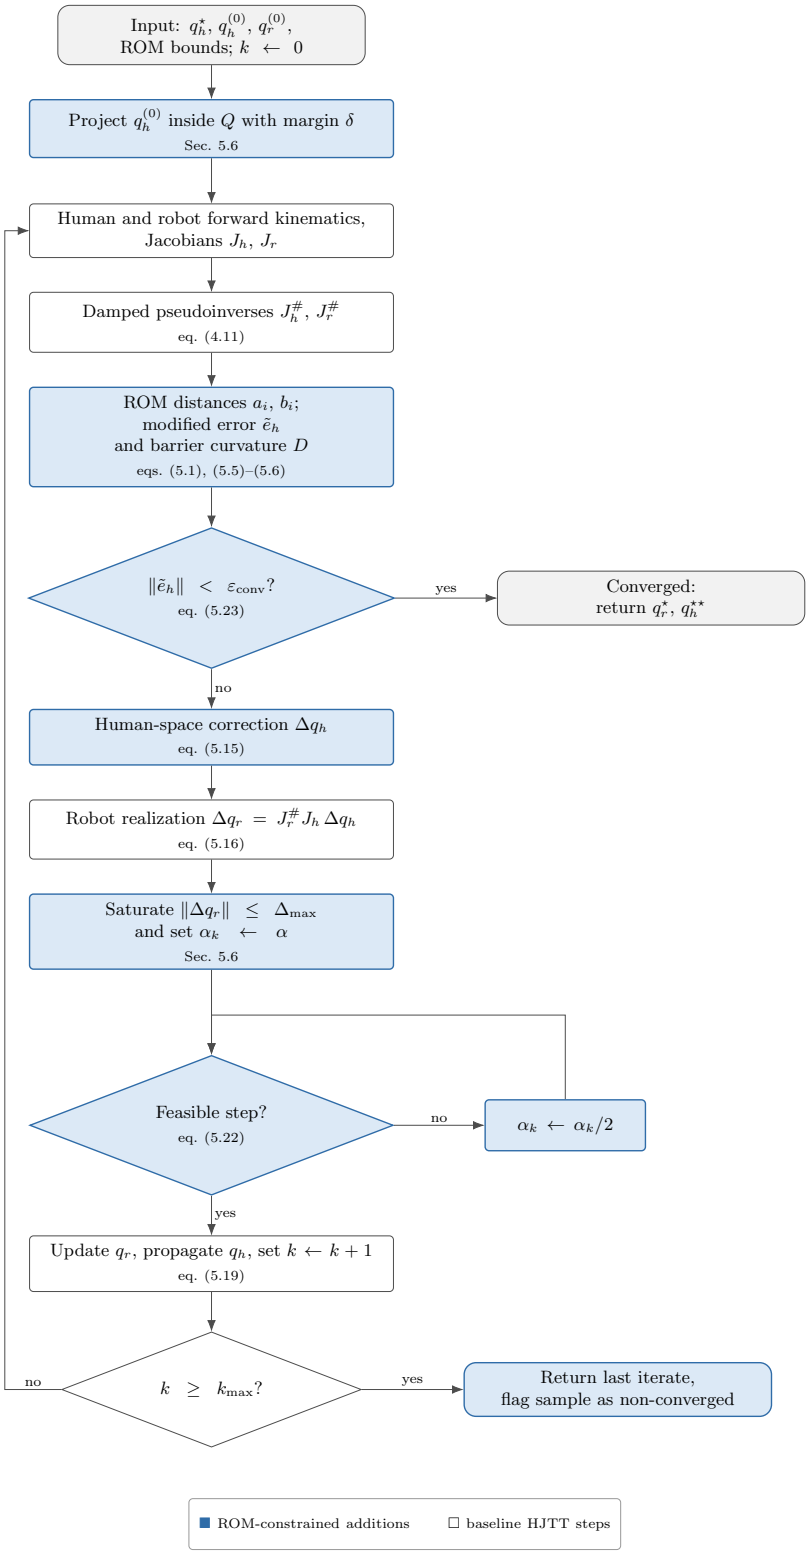


Figure 5.1: Iteration of the ROM-constrained formulation for a single sample of the reference trajectory. Blocks in blue are those introduced in this work; the remaining ones belong to the baseline HJTT algorithm.

5.5 Multiple Constrained Joints and ROM Normalization

This formulation can be extended to an arbitrary number of constrained joints, since the barrier in (5.2) is the sum of independent contributions and the matrix D in (5.6) remains diagonal. However, a single weight γ does not act uniformly on joints of different amplitude.

This becomes observable when the external factor is omitted from (5.2). As mentioned in Section 5.1, the normalization inside the logarithm disappears upon differentiation, meaning that the resulting curvature $\gamma\,(1/a_i^2 + 1/b_i^2)$ is expressed in absolute units. For a joint with a narrow range of motion, the distances from the bounds are correspondingly small, so the curvature grows accordingly. This means that a common γ would constrain such a joint far more severely than one with a wide ROM.

The factor r_i^2 removes this dependence. Substituting $a_i = \bar{a}_i r_i$ and $b_i = \bar{b}_i r_i$ into (5.6),

$$d_i = r_i^2 \left(\frac{1}{a_i^2} + \frac{1}{b_i^2} \right) = \frac{1}{\bar{a}_i^2} + \frac{1}{\bar{b}_i^2}, \tag{5.20}$$

depends only on the normalized position of the joint within its own range. When the joints are at an equal relative distance from the bounds, the ratio of the curvature of the barrier to that of the tracking term is therefore identical for all of them, and a single value of γ yields a comparable constraint on each.

However, the same factor does not make the gradient range independent. From (5.5),

$$\tilde{e}_{h,i} - e_{h,i} = \gamma\, r_i^2 \left(\frac{1}{a_i} - \frac{1}{b_i} \right) = \gamma\, r_i \left(\frac{1}{\bar{a}_i} - \frac{1}{\bar{b}_i} \right), \tag{5.21}$$

which retains a factor r_i. The normalization acts on the curvature, the quantity through which the barrier enters the Gauss–Newton step by means of the inversion in (5.15). This scaling is relevant to the behavior of the algorithm.

5.6 Feasibility, Stopping Criterion and Relation with the Baseline HJTT

Feasibility. Since Φ is only defined in the interior of Q, the iteration must stay within that domain. The initial human configuration is projected inside the admissible set with a margin $\delta > 0$. This projection is applied only to the initial human estimate and is not used as a post-processing operation on subsequent iterates. From the first iteration onward, the human configuration is propagated exclusively through the differential robot-to-human mapping (5.19).

A candidate step is accepted only if the predicted configuration satisfies the following condition:

$$q_{h,\min,i} + \delta \;<\; q_{h,i}^{(k)} + \alpha_k \left(J_h^{\#} J_r \, \Delta q_r^{(k)} \right)_i \;<\; q_{h,\max,i} - \delta, \qquad i \in \mathcal{C}. \tag{5.22}$$

The margin δ prevents the barrier from being evaluated arbitrarily close to its singularity. When (5.22) is violated the step is reduced by halving and the test is repeated, so that the accepted step size is $\alpha_k = \alpha \, 2^{-m}$, where α is the nominal gain and m the number of reductions required. Since the current iterate lies strictly inside the admissible set with margin δ and the candidate configuration tends to it as α_k vanishes, a sufficient reduction satisfies (5.22) under regular conditions. The robot increment is additionally saturated whenever its norm exceeds a prescribed threshold, so that the linearization on which the local model (5.9) relies is not applied outside its range of validity. The numerical values adopted for α, δ and the saturation threshold, together with the safeguards introduced in the implementation, are reported in Chapter 6.

Stopping criterion. Due to the bias introduced in Section 5.1, the constrained joints settle at a finite distance from their bounds at the optimum and the tracking error does not vanish whereas the gradient does. The iteration is therefore terminated based on the stationarity of the cost function,

$$\left\| \tilde{e}_h^{(k)} \right\| < \varepsilon_{\text{conv}}. \tag{5.23}$$

No separate verification of the ROM limits is required at convergence. The initial configuration is admissible by construction and every accepted step satisfies (5.22), so that all the iterates remain strictly inside Q with margin δ: admissibility is an invariant of the iteration rather than a condition to be tested at its end. The criterion (4.14) of the baseline algorithm would instead be inconsistent here, since $\|e_h\|$ is not expected to reach zero whenever the reference is not attainable. A maximum number of iterations $k_{\max}$ and a threshold on the norm of the step are retained as safeguards, and the complete procedure is summarized in Algorithm 5.1.

Relation with the baseline HJTT. For $\gamma = 0$ the barrier contribution vanishes, meaning that $\tilde{e}_h = e_h$ from (5.5) and $D = 0$ from (5.6). The correction (5.15) reduces to $\Delta q_h^{(k)} = e_h^{(k)}$ and the robot update (5.16) and (5.19) becomes

$$q_r^{(k+1)} = q_r^{(k)} + \alpha_k \, J_r^{\#}\left(q_r^{(k)}\right) J_h\left(q_h^{(k)}\right) e_h^{(k)}, \tag{5.24}$$

which coincides with (4.11), while the human update coincides with (4.13) and the stopping criterion (5.23) reduces to (4.14). For $\gamma = 0$ the correction law therefore reduces to that of the baseline HJTT; if, in addition, the feasibility safeguards remain inactive and the nominal full step is accepted, the complete iteration coincides with the baseline formulation. For feasible references sufficiently far from the ROM boundaries, the influence of the barrier remains limited.

Algorithm 5.1 ROM-constrained HJTT optimization procedure for a given desired configuration $q_h^\star$

Require: $q_h^\star$, $q_h^{(0)}$, $q_r^{(0)} \in \mathbb{R}^{n_r}$, H_{DH}, R_{DH}, $\mathcal{C}$, $q_{h,\min}$, $q_{h,\max}$, $\gamma \geq 0$, $\alpha \in (0,1]$, $\lambda > 0$, $\delta > 0$, $\Delta_{\max} > 0$, $\varepsilon_{\text{conv}} > 0$, $k_{\max}$
Ensure: $q_r^\star$, $q_h^{\star\star}$, δ_{conv}

$k \leftarrow 0$, $\delta_{\text{conv}} \leftarrow$ FALSE
project $q_h^{(0)}$ into the interior of Q with margin δ
while $\delta_{\text{conv}} =$ FALSE **and** $k < k_{\max}$ **do**
    compute the human and robot forward kinematics and assemble $J_h(q_h^{(k)})$, $J_r(q_r^{(k)})$
    $J_h^\#$, $J_r^\#$ as in (4.12)
    $a_i \leftarrow q_{h,i}^{(k)} - q_{h,\min,i}$, $b_i \leftarrow q_{h,\max,i} - q_{h,i}^{(k)}$, $i \in \mathcal{C}$ ▷ eq. (5.1)
    $\tilde{e}_h^{(k)} \leftarrow$ modified error, $D^{(k)} \leftarrow$ barrier curvature ▷ eqs. (5.5), (5.6)
    if $\|\tilde{e}_h^{(k)}\| < \varepsilon_{\text{conv}}$ **then**
        $\delta_{\text{conv}} \leftarrow$ TRUE; **break** ▷ stationarity, eq. (5.23)
    end if
    $\Delta q_h^{(k)} \leftarrow (I + \gamma D^{(k)})^{-1} \tilde{e}_h^{(k)}$ ▷ GGN step, eq. (5.15)
    $\Delta q_r^{(k)} \leftarrow J_r^\# J_h \, \Delta q_h^{(k)}$ ▷ robotic realization, eq. (5.16)
    if $\|\Delta q_r^{(k)}\| > \Delta_{\max}$ **then**
        $\Delta q_r^{(k)} \leftarrow \Delta_{\max} \, \Delta q_r^{(k)} / \|\Delta q_r^{(k)}\|$ ▷ step saturation
    end if
    $\alpha_k \leftarrow \alpha$; **while** (5.22) is violated **do** $\alpha_k \leftarrow \alpha_k / 2$ ▷ feasibility, eq. (5.22)
    $q_r^{(k+1)} \leftarrow q_r^{(k)} + \alpha_k \Delta q_r^{(k)}$ ▷ eq. (5.19)
    $q_h^{(k+1)} \leftarrow q_h^{(k)} + J_h^\# J_r (q_r^{(k+1)} - q_r^{(k)})$ ▷ eq. (5.19)
    $k \leftarrow k + 1$
end while
$q_r^\star \leftarrow q_r^{(k)}$, $q_h^{\star\star} \leftarrow q_h^{(k)}$

5.7 Application to a complete trajectory

As in the baseline algorithm, the procedure described above refers to a single desired configuration and is applied sequentially to all the samples of the recorded trajectory $\{q_h^\star(t_i)\}_{i=1}^N$. The ROM limits are determined once for the subject and held constant along the trajectory, so that the admissible set Q does not depend on the sample being processed; the procedure adopted to determine them is described in Chapter 6.

The initialization of the first sample follows (4.15): the robot configuration is obtained by solving the inverse kinematics for the position given by the forward kinematics of the arm model in the first desired configuration and for the prescribed end-effector orientation. The initial human estimate is then projected inside Q as described in Section 5.6.

For the subsequent samples the initialization differs from that of the baseline. In Section 4.2.5 the human configuration is initialized from the previous sample of the reference trajectory, which is legitimate there because the baseline drives the human configuration towards the reference and the two coincide at convergence. Under the constrained formulation this is no longer the case: the minimizer of (5.2) is displaced towards the interior of the admissible set, and whenever the reference is not attainable the two differ by construction. The human configuration is therefore initialized from the configuration obtained at the previous sample,

$$q_r^{(0)}(t_i) = q_r^\star(t_{i-1}), \qquad q_h^{(0)}(t_i) = q_h^{\star\star}(t_{i-1}), \qquad i > 0,$$

so that both initial guesses are inherited from the previously computed pair, promoting continuity between consecutive trajectory samples.

When the stopping criterion (5.23) is not met within $k_{\max}$ iterations, the last iterate is retained as the solution for that sample and its index is recorded, so that the number and the position of the non-converged samples are available for the analysis of the results. Since every accepted step is required by (5.22) to keep the constrained joints inside Q, the retained configuration is admissible even when the stationarity threshold has not been reached.

The final result is a robot trajectory $\{q_r^\star(t_i)\}_{i=1}^N$ designed to balance tracking of the recorded reference with separation from the prescribed ROM boundaries, while keeping the predicted constrained human joint coordinates within their admissible ranges.

Chapter 6

Experimental Setup

6.1 Overview and objectives

The constrained formulation introduced in Chapter 5 was evaluated through experimental tests on the real robotic platform. The experiments were designed to assess whether the imposed range-of-motion limits are effectively enforced on the joints of the guided arm, while the overall movement demonstrated by the therapist is still reproduced.

Differently from [4], the validation was carried out entirely in the experimental setting, without a preliminary simulation phase. This chapter describes the subject, the acquisition procedure, the definition of the limits and the parameters adopted, while the results are presented in Chapter 7.

6.1.1 Subject description

A single healthy subject took part in the experiments, acting both as the therapist, by recording the reference movements, and as the guided subject, by holding the robot end-effector during the executions. The subject participated voluntarily and was free from any condition affecting the mobility of the upper limb; no patients were involved.

The choice of a single subject covering both roles has a methodological consequence that is worth making explicit. Since the reference and the guided arm share the same anthropometry, the adaptation to a different limb morphology, which is the contribution of the baseline framework and was validated in [5], is inactive here: the kinematic chain of the recording phase and the one of the execution phase coincide. The experiments therefore isolate the effect of the range-of-motion constraint, which is the object of this work, from the morphological adaptation that the method inherits. The residual discrepancy between predicted and measured configurations reported in Chapter 7 is consequently attributable to the kinematic model itself, and not to a difference in segment lengths between demonstrator and patient.

The structure of the right arm was modeled using the Denavit–Hartenberg parameters reported in Table 4.1, with the link lengths set to the measured values of the subject. Table 6.1 summarizes the anthropometric parameters used in the calibration of the inertial suit and in the algorithm.

Table 6.1: Anthropometric parameters of the subject. The shoulder height is measured from the floor with the subject seated, in the fixed chair position adopted throughout the experiments.

Parameter	Value
Upper arm length l_{UA} [m]	0.30
Forearm length l_{FA} [m]	0.23
Seated shoulder height [m]	1.00

6.1.2 Acquisition and reference trajectories

Before each session the inertial suit was calibrated with the procedure described in Section 3.2, and the recorded joint angles were read and converted as described in Sections 3.2.1 and 4.1. The system records at 60 Hz.

Two movements were recorded. In both of them the shoulder flexion–extension joint carries the largest excursion, while the second most solicited joint differs, so that the behavior of the constrained formulation could be verified without the outcome depending on which particular joint is being limited. The excursions measured on the two references are reported in Table 6.2.

Table 6.2: Joint excursions of the two reference movements, measured on the recorded trajectories.

	SFE	SAA	SIER	EFE
Movement 1 [deg]	87.7	28.5	62.8	21.8
Movement 2 [deg]	68.0	55.5	62.0	17.5

The first movement is dominated by shoulder flexion, with a secondary contribution of the internal–external rotation. The second combines flexion with a larger abduction component. Both were recorded by the subject in free space while wearing the sensors.

The two references are shown in Figure 6.1. In both of them the flexion–extension of the shoulder describes a single excursion and return, and the elbow stays within a narrow band, so that the movement is essentially organized around the shoulder. What distinguishes the two cases is the joint that accompanies the flexion: the internal–external rotation in the first, which follows the same profile with a smaller amplitude, and the abduction–adduction in the second, which rises earlier than the flexion and holds a plateau while the flexion reaches its own maximum. The two references last 10.3 s and 9.7 s, corresponding to 618 and 583 samples at the acquisition rate of 60 Hz.

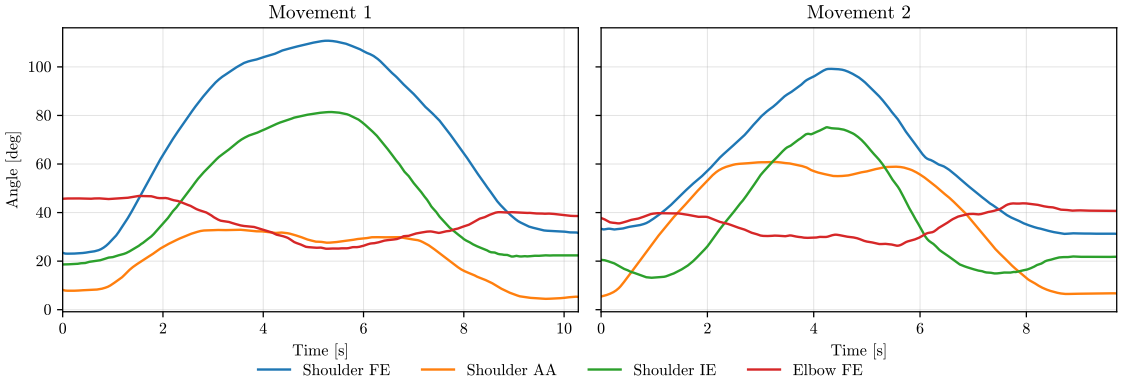


Figure 6.1: Joint trajectories of the two reference movements, recorded from the therapist in free space. Left: Movement 1. Right: Movement 2. The excursions of these trajectories are the ones reported in Table 6.2.

6.2 Definition of the range-of-motion limits

The limits were not assigned as fixed anatomical values but derived from the range of the recorded reference, so that they fall by construction within the movement demonstrated by the therapist. For each constrained joint i, denoting by $q_i^{\min\star}$ and $q_i^{\max\star}$ the extreme values of the reference and by $\Delta_i^\star = q_i^{\max\star} - q_i^{\min\star}$ its range, the limits that define the admissible set (4.15) are

$$q_i^{\min} = q_i^{\min\star} + \rho_{\text{low}}\,\Delta_i^\star, \qquad q_i^{\max} = q_i^{\min\star} + \rho_{\text{high}}\,\Delta_i^\star, \tag{6.1}$$

with $\rho_{\text{low}} = 0$ and $\rho_{\text{high}} = 0.80$ in the tests reported here. The lower limit therefore coincides with the minimum of the reference, while the upper limit removes the last 20% of its excursion. For the abduction–adduction joint the two fractions are mirrored with respect to the maximum of the reference, because of the sign convention adopted for that angle.

The resulting bounds are reported in Table 6.3. The upper limit removes between 11.1° and 17.5° from the excursion of the reference, so that in every case the demonstrated movement enters a region that the constrained formulation must avoid.

Table 6.3: Range-of-motion limits imposed in the two movements, obtained from (6.1) with $\rho_{\text{low}} = 0$ and $\rho_{\text{high}} = 0.80$. All values in degrees.

Movement	Joint	$q_i^{\min\star}$	$q_i^{\max\star}$	$\Delta_i^\star$	$q_i^{\min}$	$q_i^{\max}$
Movement 1	Shoulder FE	23.1	110.8	87.7	23.1	93.2
	Shoulder IE	18.6	81.4	62.8	18.6	68.9
Movement 2	Shoulder FE	31.2	99.2	68.0	31.2	85.6
	Shoulder AA	5.4	60.9	55.5	5.4	49.8

In each movement two joints were constrained, chosen so that their excursions in the reference are comparable and so that the second constrained joint differs between

the two cases: shoulder flexion–extension and internal–external rotation in the first movement, shoulder flexion–extension and abduction–adduction in the second. The requirement of comparability is not arbitrary: constraining simultaneously two joints whose excursions differ by a large factor makes the feasibility test of Section 6.3 reject every candidate step, since the joint with the smaller range is already close to its own limits for any displacement that is meaningful for the other. This is a condition of applicability of the method rather than a limitation of its formulation.

It is important to state explicitly that the limits obtained in this way are *not* physiological limits. They are derived from a recorded movement and serve to verify that the algorithm enforces a prescribed bound; in a clinical application the bounds would be provided by the assessment of the patient, and the procedure described here would simply take them as input.

6.3 Algorithm parameters and numerical safeguards

Table 6.4 reports the parameters used in all the tests. The same values were adopted for both movements, so that the comparison is not affected by a re-tuning between cases.

Table 6.4: Parameters of the constrained algorithm.

Parameter	Symbol	Value
Barrier weight	γ	10^{-2}
Nominal step gain	α	0.05
Pseudoinverse damping	λ	0.02
Safety margin from the limits	δ	0.5°
Maximum robot step	$\Delta_{\max}$	5.0°
Convergence threshold on $\|\tilde{e}_h\|$	$\varepsilon_{\text{conv}}$	10^{-3}
Maximum iterations per sample	$k_{\max}$	4000
Step-norm threshold	ε_Δ	10^{-8}
Admitted residual ROM violation	—	1.0°
Maximum number of step reductions	—	20

The first seven entries of Table 6.4 are the numerical parameters appearing in Algorithm 5.1. The remaining three quantify safeguards of the implementation that are not parameterized there: the step-norm threshold, mentioned in Section 5.6 but not among the inputs of the algorithm, and the two provisions described below, which belong to the implementation alone. The backtracking of Section 5.6 is limited to twenty reductions, after which the step is applied with the last value of α_k; as a consequence, the stopping criterion of the implementation verifies, in addition to the stationarity condition (5.23), that the residual violation of the ROM limits remains below one degree. Both measures protect against the pathological case in which no admissible step is found, which the theoretical argument of Section 5.6 excludes under regular conditions.

Two further provisions complete the numerical treatment. The distances a_i and b_i entering the barrier are clamped from below before being inverted, so that the

logarithmic terms remain finite even if an iterate were to reach a bound. The iteration is additionally stopped when the norm of the accepted step becomes negligible, which terminates the procedure when no further progress is possible.

The barrier is evaluated with a fixed weight throughout, without any continuation scheme on γ, for the reasons discussed in Section 5.6.

6.4 Experimental implementation

The experiments were performed in the Intelligent Robotics Laboratory, with the seven-degree-of-freedom collaborative robot described in Section 3.1. The robot was mounted on a workbench with its base 0.90 m above the floor, and the subject was seated approximately 0.90 m from the robot base, as shown in Figure 6.2. The hand was coupled to the end-effector through the dedicated 3D-printed handle shown in Figure 3.2, which guarantees a stable and repeatable connection during the execution.



Figure 6.2: Experimental setup in the Intelligent Robotics Laboratory. The robot base is 0.90 m above the floor and the seat of the subject is placed at 0.90 m from it.

The procedure was the following. The subject, wearing the inertial suit, first

performed the movement in free space, producing the reference trajectory. The reference was then processed offline through both the original HJTT formulation and the constrained one, obtaining two robot joint trajectories. The subject, still wearing the sensors, then grasped the handle and let the robot guide the arm along the computed trajectories, so that the joint angles actually realized were recorded rather than inferred from the robot configuration.

Each constrained trajectory was executed ten times, following the procedure adopted in [4], so that mean trajectories and variability bands could be computed across repetitions. For the unconstrained trajectory, a single execution per movement was retained as the term of comparison. Figure 6.3 shows three instants of one constrained execution for each of the two movements.

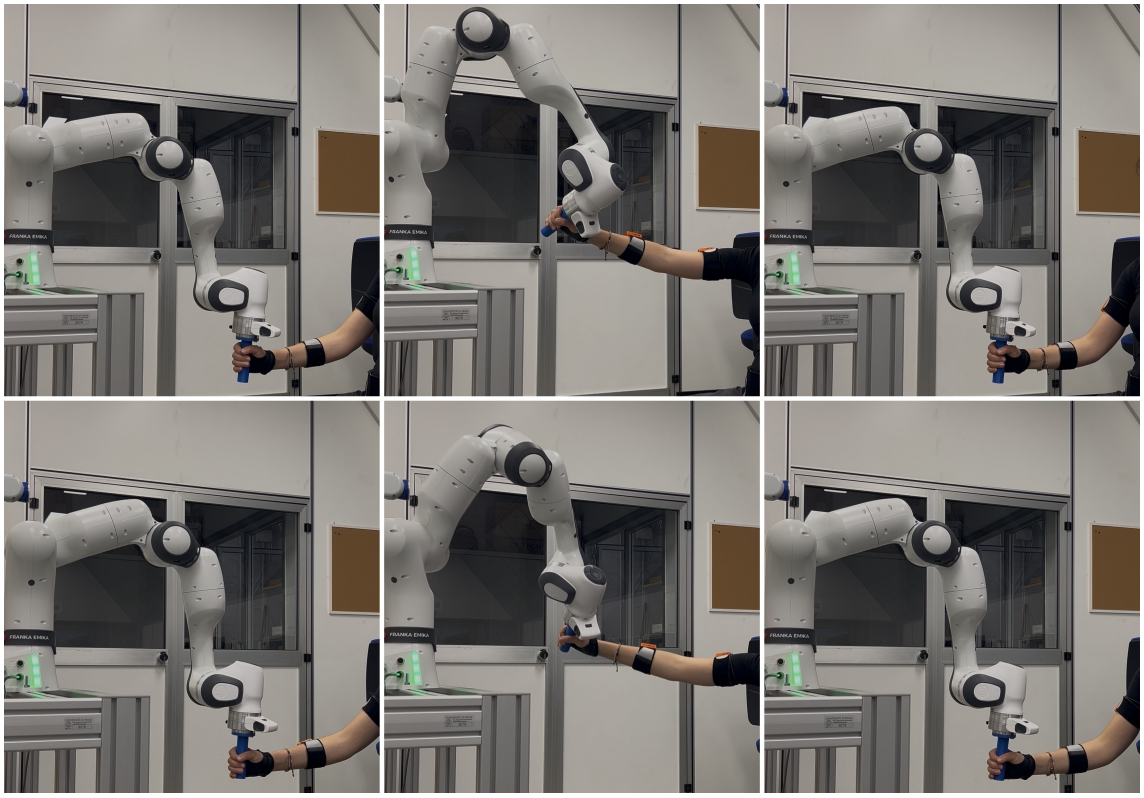


Figure 6.3: Initial, intermediate and final instants of the guided execution of the *constrained* trajectory, for the two movements. Top row: Movement 1. Bottom row: Movement 2. The intermediate frame of each row corresponds to the peak of the dominant constrained joint, the shoulder flexion–extension in the first movement and the abduction–adduction in the second.

The recording of the movement and the start of the robot execution were commanded from two separate workstations and could not be synchronized exactly. Each acquisition therefore begins with a preamble of variable duration that does not belong to the movement; the consequence of this on the processing of the repetitions is addressed in Chapter 7.

6.5 Software integration and experimental workflow

The experiments were carried out within the software framework developed in [15] and extended in [4], the same framework later integrated into the platform of [16], which relies on ROS 2 for the communication with the robot and provides the graphical interface used by the clinician to record a movement, load a previously acquired trajectory and execute it on the robot. For this work the framework was used without modifications to its structure: the constrained trajectories were computed offline and loaded through the same interface used for the unconstrained ones.

In operational terms the session is organized in three stages, summarized in Figure 6.4. In the recording stage the subject, wearing the inertial suit, performs the movement in free space while MVN Analyze acquires the segment orientations; the joint angles are then exported and converted according to the conventions of Section 3.2.1. In the generation stage the reference is processed offline, once through the baseline formulation and once through the constrained one, producing two robot joint trajectories for the same demonstrated movement. In the execution stage the trajectory is loaded through the clinician interface and played on the manipulator while the subject, still instrumented, holds the handle; the joint angles realized during the guided movement are recorded at the same time by the same inertial system, so that the quantity used for the verification is measured on the arm rather than inferred from the robot configuration. Execution and acquisition are therefore simultaneous, and not two successive steps.

Keeping the generation offline is a deliberate choice at this stage of the work: it separates the verification of the formulation from any constraint on computation time, and it allows the same trajectory to be replayed identically across repetitions, which is what makes the variability bands of Chapter 7 meaningful.

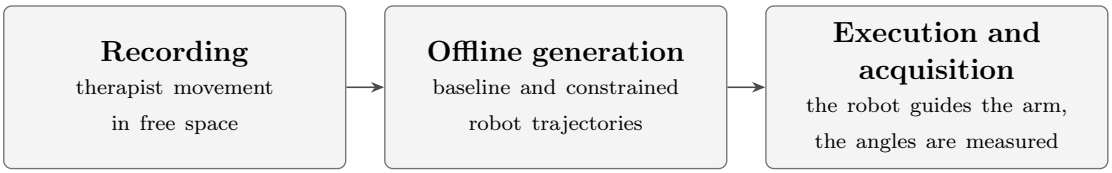


Figure 6.4: Schematic representation of the experimental workflow.

Chapter 7

Results

7.1 Introduction to results

For each of the two movements, the robot trajectory was computed twice: once with the original HJTT formulation and once with the constrained formulation of Chapter 5. Each constrained trajectory was executed ten times, so that a mean joint trajectory and a min–max interval could be computed across repetitions, while a single unconstrained execution per movement was retained as a term of comparison. During every execution the subject held the end-effector while wearing the inertial suit, so that the joint angles reported here were measured on the arm rather than inferred from the robot configuration.

Because the motion-capture acquisition and the robot execution were started independently, the repetitions exhibited different initial and final offsets. The constrained trials were therefore aligned using the rising edge of the SFE trajectory before computing the mean and the min–max interval.

The generation of the constrained trajectories converged for every reference sample, with 0 out of 618 samples flagged as non-converged for Movement 1 and 0 out of 583 for Movement 2. The limits imposed in the two movements are those reported in Section 6.2.

7.2 Predicted human configuration

Before examining the measured trajectories, it is useful to look at the configuration that the algorithm predicts while computing the robot trajectory. The formulation operates on the human joint configuration reconstructed through the kinematic model, and it is on that configuration that the barrier terms act.

Figures 7.1 and 7.2 report, for the two movements, the predicted human joint trajectory superimposed on the therapist reference, with the imposed limits shown on the constrained joints. On the constrained joints the predicted trajectory separates from the reference before the latter reaches the boundary and remains inside the admissible set throughout the movement, reaching 90.9° and 67.2° in the first movement, against limits of 93.2° and 68.9°, and 83.8° and 48.3° in the second, against 85.6° and 49.8°. On the joints that are not constrained the two curves coincide, since for those joints the correction reduces to that of the original formulation.

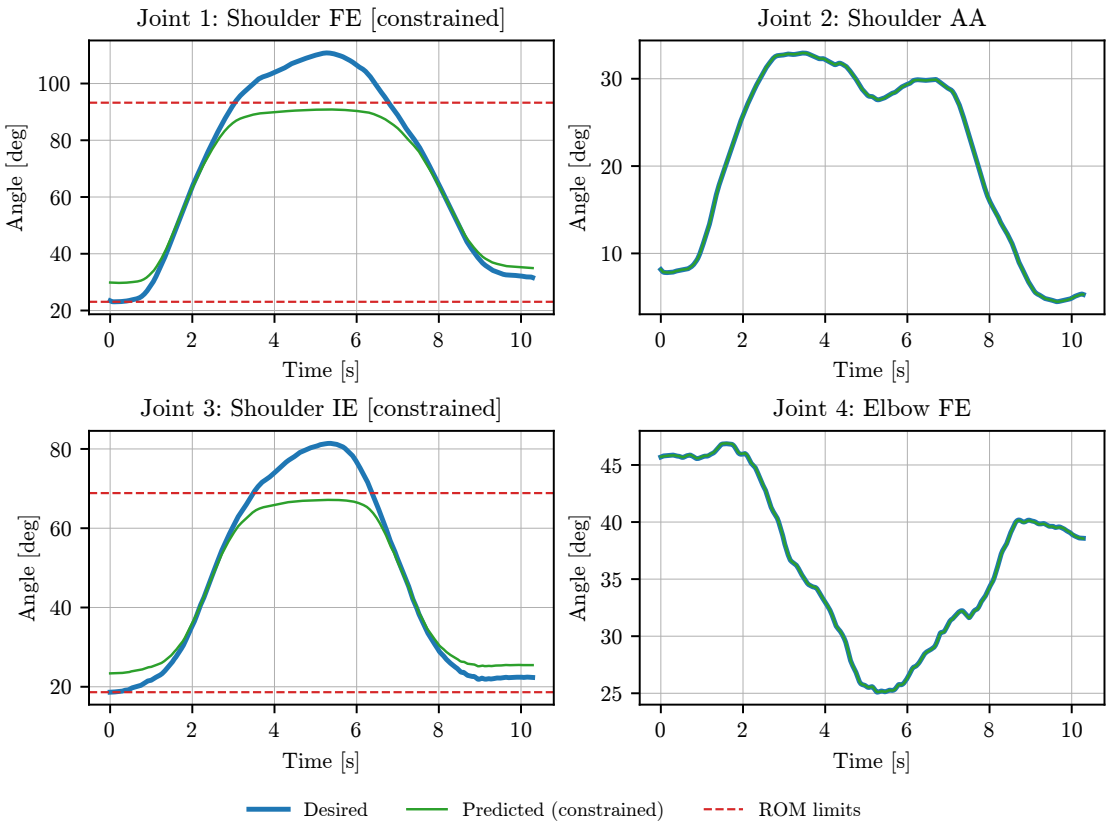


Figure 7.1: Movement 1: human joint configuration predicted during trajectory generation (solid green), superimposed on the therapist reference (solid blue). Dashed lines indicate the imposed ROM limits on the constrained joints.

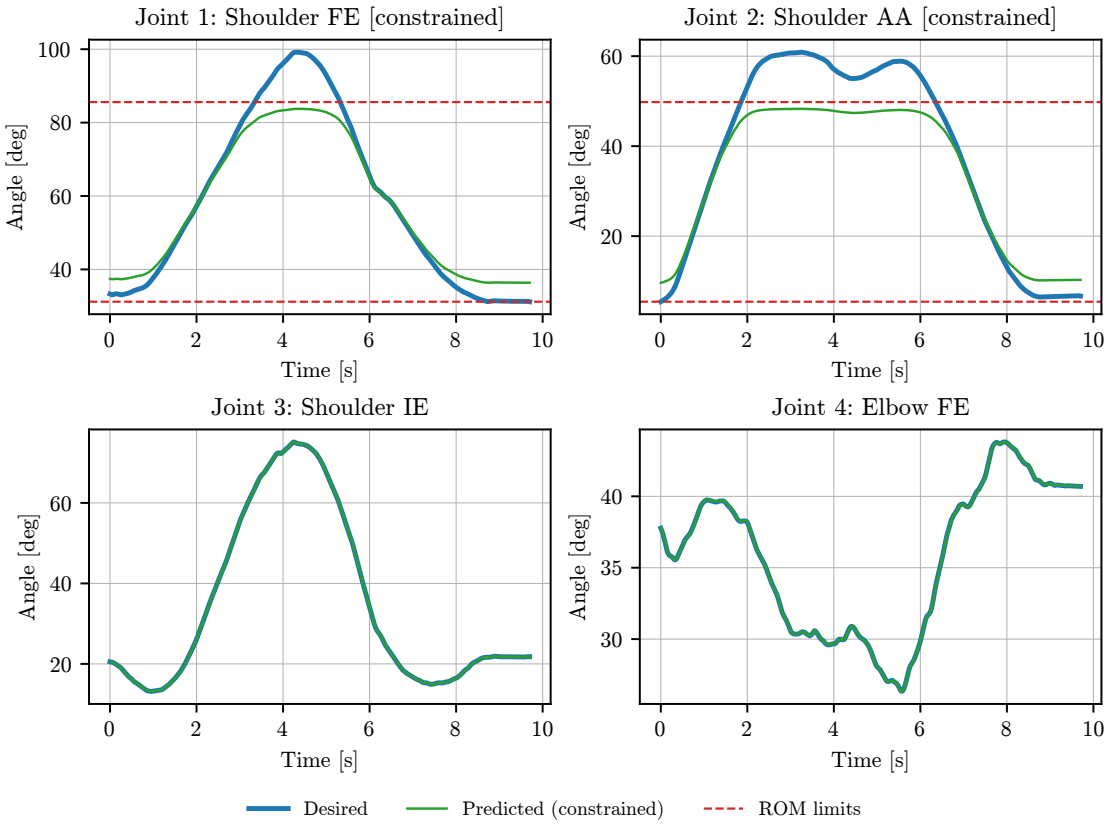


Figure 7.2: Movement 2: predicted human joint configuration, with the same representation adopted in Figure 7.1.

A further property of these trajectories is apparent from the figures. The predicted motion remains smooth where the constraint becomes active: the curve bends away from the reference gradually, well before the boundary is reached, and rejoins it just as gradually once the reference returns inside the admissible set. This is a consequence of the formulation rather than an incidental feature. The barrier contributes over the whole admissible domain and its influence grows continuously as the boundary is approached, so the correction is distributed over a portion of the movement instead of being concentrated at the instants in which the limit would be crossed. Clipping the unconstrained HJTT trajectory at the prescribed bounds would instead produce a signal with discontinuous derivative at the onset and release of saturation, and a flat segment in between. The distinction matters in this application, because the resulting trajectory is executed by a robot that guides the arm of a patient: a motion that varies smoothly is more appropriate for that purpose than one obtained by truncating the reference. It should also be recalled, from Section 4.4, that clipping would in any case break the kinematic consistency between the two chains, since the saturated human configuration would no longer correspond to the computed robot trajectory.

These are the values reported in the last column of Tables 7.1 and 7.2.

The constraint is formulated in the human joint space, but since the two chains are rigidly coupled it also modifies the Cartesian path travelled by the end-effector, and therefore by the hand. Figure 7.3 compares the two generated trajectories in the workspace of the manipulator. The reduction imposed on the joint excursions appears there as a shorter path: the vertical span decreases from 52.6 to 36.4 cm in the first movement and from 41.8 to 36.4 cm in the second, and the overall path length by 22% and 13% respectively.

The constrained path is not a scaled copy of the unconstrained one. The span along y increases in both cases, from 14.7 to 18.2 cm and from 7.5 to 11.7 cm, which follows from the correction being defined in the human joint space: the Cartesian path is the one produced by the differential coupling and is not itself the object of the optimization. The mean distance between the two paths, evaluated sample by sample, is 16.2 cm in the first movement and 15.8 cm in the second.

All the quantities reported so far describe what the algorithm expects the arm to do; the sections that follow report what the arm actually did.

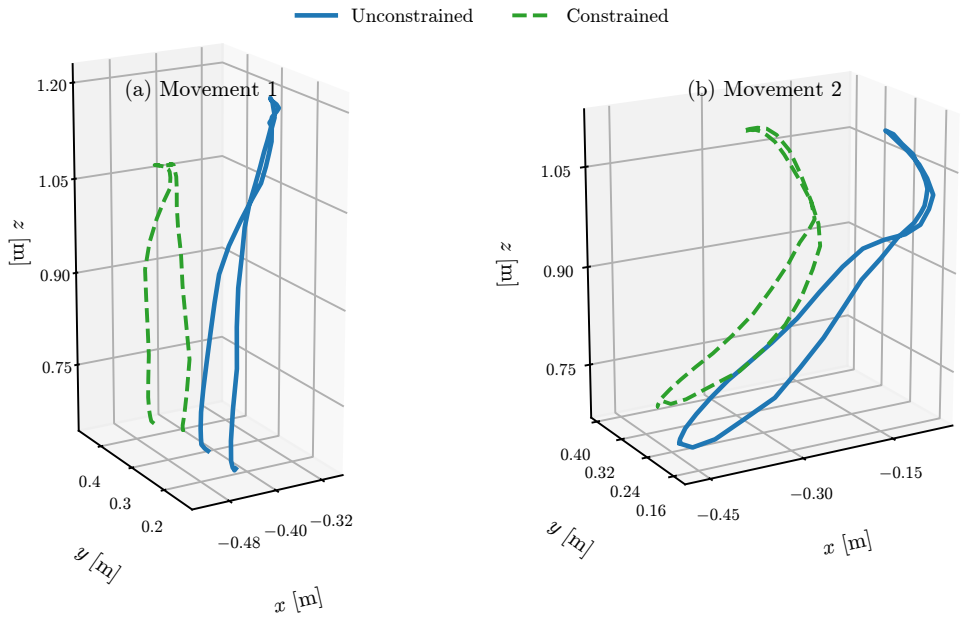


Figure 7.3: End-effector trajectory generated with and without the ROM constraint, for the two movements. The three axes are drawn to the same scale.

7.3 Experimental results

7.3.1 Movement 1

Figure 7.4 reports the measured joint trajectories for the first movement. For each joint the solid green line is the mean over the ten constrained repetitions and the shaded region is the min–max interval across trials; the therapist reference and the unconstrained execution are superimposed, and the imposed limits are shown as dashed lines on the two constrained joints.

The constrained joints, shoulder flexion–extension and internal–external rotation, exhibit a clear reduction with respect to the reference. Consistently with the finite-weight formulation of Section 5.1, the constrained mean departs from the reference before the latter reaches the imposed limit, and the separation increases as the boundary is approached. None of the ten repetitions exceeds the upper limit at any instant of the tracking window. The variability across repetitions is limited: after alignment the mean band width is 4.8° on SFE and 5.3° on SIER, and remains below 9° on the two unconstrained joints.

The last column reports the configuration predicted by the kinematic model during trajectory generation. The predicted constrained trajectory reached 90.9° on SFE and 67.2° on SIER, remaining inside the admissible set at a distance of 2.3° and 1.7° from the respective boundaries, whereas the corresponding measured mean peaks were 72.4° and 53.2°. This difference highlights that the kinematic model predicts the human response only approximately, which motivates the experimental verification introduced in Section 6.1.

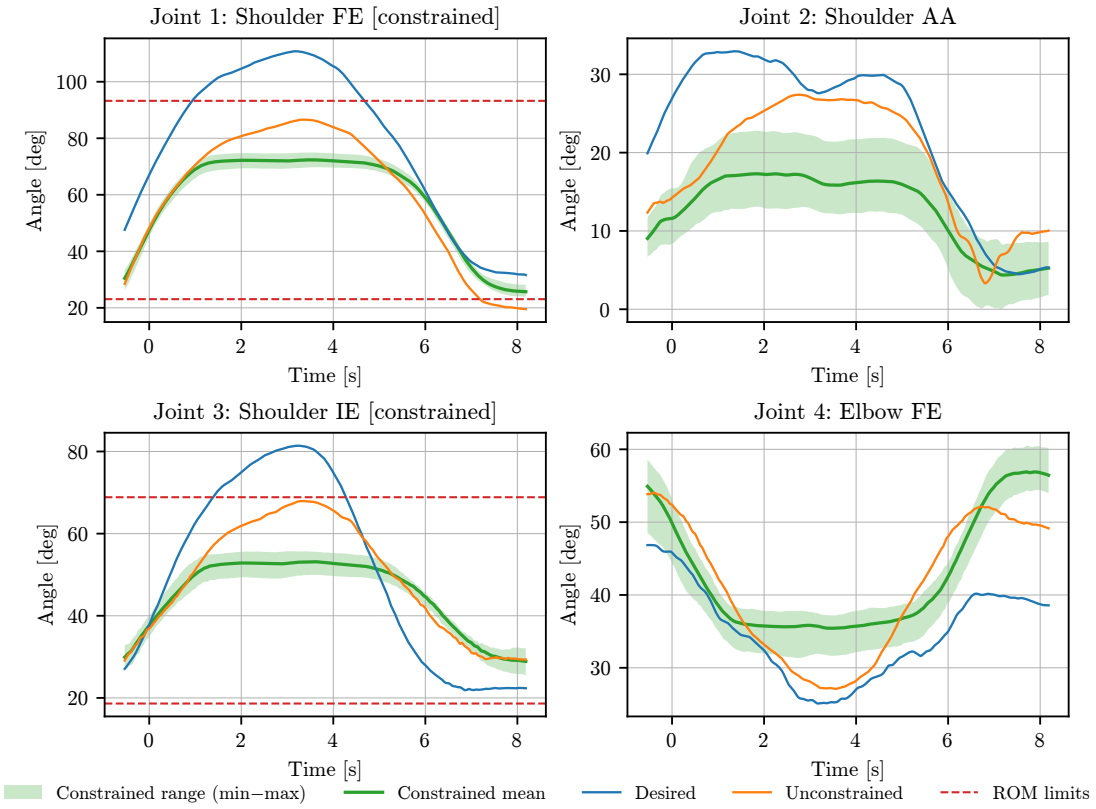


Figure 7.4: Movement 1: measured human joint trajectories. Mean over the ten constrained repetitions (solid green), min–max interval across trials (shaded), therapist reference and unconstrained execution. Dashed lines indicate the imposed ROM limits on the constrained joints.

Table 7.1: Movement 1: peak values of the constrained joints [deg]. Angles are reported according to the anatomical sign convention adopted in the experimental plots.

Joint	ROM limit	Reference	Unconstrained	Measured mean peak	Predicted
Shoulder FE	93.2	110.8	86.6	72.4	90.9
Shoulder IE	68.9	81.4	68.0	53.2	67.2

7.3.2 Movement 2

Figure 7.5 reports the measured joint trajectories for the second movement, in which the constrained joints are shoulder flexion–extension and abduction–adduction. The same reduction with respect to the reference is observed on both constrained joints, and again no repetition exceeds the upper limit during the tracking window. After alignment the mean band width is 4.8° on SFE and 6.3° on SAA, and below 5° on the two unconstrained joints.

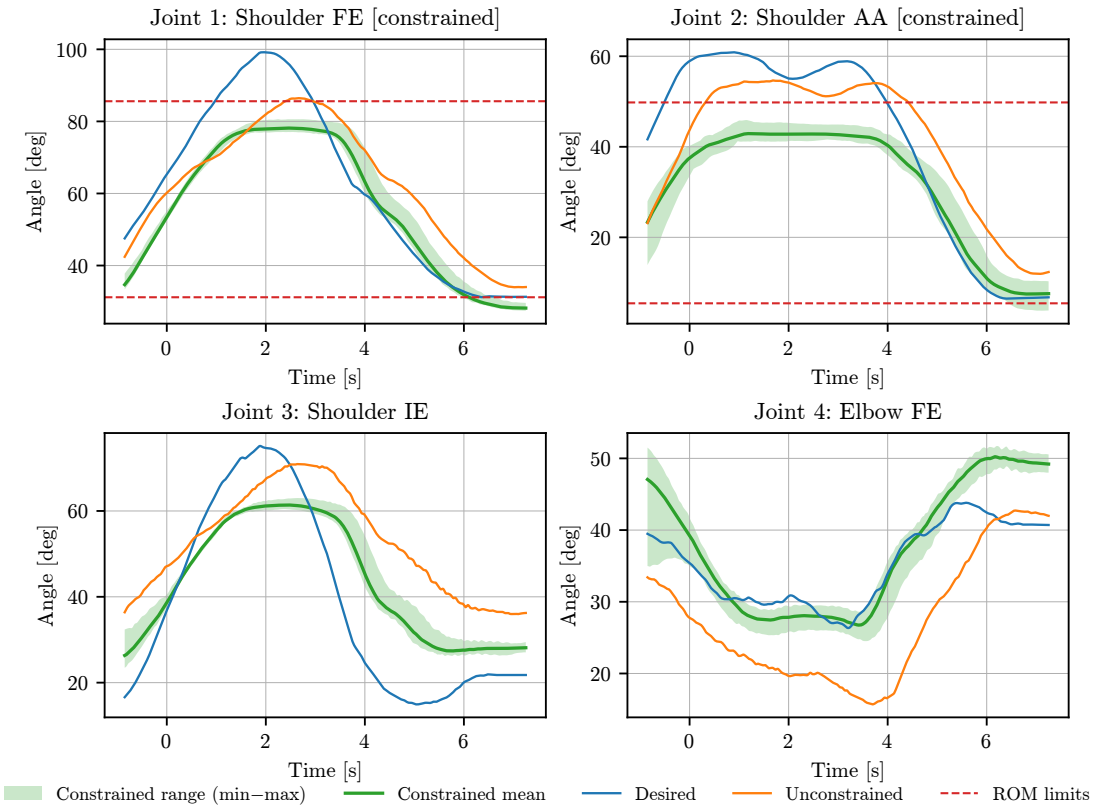


Figure 7.5: Movement 2: measured human joint trajectories, with the same representation adopted in Figure 7.4.

Table 7.2: Movement 2: peak values of the constrained joints [deg]. The SAA values are reported with the anatomical sign, obtained from the kinematic coordinate through the sign transformation of Section 4.1.

Joint	ROM limit	Reference	Unconstrained	Measured mean peak	Predicted
Shoulder FE	85.6	99.2	86.5	78.1	83.8
Shoulder AA	49.8	60.9	54.7	42.9	48.3

Here the predicted configuration reached 83.8° on SFE and 48.3° on SAA, against measured mean peaks of 78.1° and 42.9°. The agreement between prediction and

measurement is therefore closer than in Movement 1, where the same comparison differed by approximately eighteen degrees on SFE.

7.4 Overall observations

Compliance with the imposed limits was verified separately for each trial, on the raw recorded signal and therefore independently of the alignment adopted for the figures, restricting the analysis to the window in which the trajectory is actually being executed. All the 20 constrained executions remain within the imposed limits during the movement, and the upper limit, which is the one the constraint is intended to enforce, is never exceeded for any value of the margin used to define the window, up to one second at both ends.

Outside that window the picture is different, and for a reason that follows from the definition of the limits. With $\rho_{\text{low}} = 0$ the lower bound coincides by construction with the minimum of the reference, which lies above the resting posture of the arm: lower-bound crossings therefore appear in the phases in which the robot brings the arm to the starting configuration or returns it to rest. Over the complete recordings a single upper-bound crossing was also observed, of approximately 1.6° on the SAA signal of one trial, after the guided trajectory had ended.

The behavior of the retained unconstrained HJTT execution differed between the two movements. In Movement 1 the measured unconstrained peaks remained inside both prescribed bounds, at 86.6° against a limit of 93.2° on SFE and 68.0° against 68.9° on SIER. In Movement 2 both bounds were exceeded, by approximately five degrees on SAA and by less than one degree on SFE. The difference originates in the discrepancy between the predicted and the measured configuration rather than in the trajectory generation. On the reference of Movement 1 the baseline formulation reproduces the demonstrated joint trajectory to within its convergence tolerance, so the predicted unconstrained peak coincides with the reference value of 110.8°; the measured one is 86.6°, that is 24.2° lower. Since the limit imposes a restriction of 17.5° with respect to the reference, and the measured execution already falls further below it than that, the bound leaves nothing to remove on the measured signal. In Movement 2 the same gap is 12.7°, roughly half, and the unconstrained execution does cross both prescribed bounds.

This does not mean that the constraint was inactive in Movement 1, since the constrained execution reached 72.4° against the 86.6° of the unconstrained one: it means that Movement 1 did not provide a strong experimental stress test of the constraint, while in Movement 2 the two formulations are compared under a condition in which the limits are actually binding.

Taken together, the two movements show the same behavior under two different constraint configurations: the prescribed limits are respected throughout the guided motion, the repetitions show limited variability once the acquisitions are aligned, and the measured peaks of the constrained joints remain farther from the prescribed upper boundaries than the corresponding values predicted by the kinematic model, not only for the mean trajectories but also in every individual repetition.

Chapter 8

Conclusion

8.1 Conclusions

The objective of this thesis was to address a limitation that emerges when the Human Joint Trajectory Tracking framework is applied to a patient whose joint mobility is restricted. The framework is designed to minimize the discrepancy between the joint motion demonstrated by the therapist and the one induced in the patient; when the range of motion available to the patient is reduced with respect to that of the therapist, more accurate tracking may drive the predicted human configuration towards excursions that the demonstrator could perform but the patient may not. Accuracy of tracking and compatibility with the available range of motion are, in that situation, in tension, and in a passive exercise the patient cannot resolve it by opposing the movement.

The approach proposed here addresses that tension within the generation of the trajectory rather than at the level of the low-level controller. The tracking objective of the original formulation is augmented with logarithmic barrier terms associated with the prescribed joint limits, and the resulting problem is solved by a generalized Gauss–Newton scheme. The correction prescribed by the cost function is computed in the human joint space, where the Hessian is diagonal and positive definite, and is then realized on the manipulator through the same damped differential coupling adopted by the baseline algorithm. This avoids introducing an additional robot-space norm criterion. In the undamped full-rank limit, the resulting increment preserves the human–robot differential coupling exactly; with the damped inverses used in the implementation, it provides a regularized realization of the same coupling structure. The barrier terms are scaled by the width of each admissible range so that, at equal normalized positions within their ROMs, their curvature contribution is independent of the absolute range width. When $\gamma = 0$, the correction law reduces to that of the original HJTT formulation.

The method was evaluated experimentally on a collaborative manipulator, on two movements with different kinematic content and two different pairs of constrained joints, with ten repetitions of each constrained trajectory. In all the twenty executions the constrained joints remained within the prescribed limits throughout the guided movement, and the upper limit, which is the bound the constraint is intended to enforce, was never exceeded. The trajectory generation converged for

every sample of both references. Once the acquisitions were synchronized, the dispersion across repetitions was limited, with a mean min–max band of 4.8° on the joint carrying the largest excursion in both movements.

The experiments also produced a result that a simulated validation could not have provided. The formulation acts on the human joint configuration predicted by the kinematic model, and the comparison between that prediction and the motion measured on the subject revealed a discrepancy that is not negligible: in the first movement the predicted flexion–extension peak exceeded the measured one by approximately eighteen degrees, while in the second the two agreed within six. In every constrained joint of both movements the measured peak remained farther from the imposed boundary than the predicted one, in each individual repetition and not only on average. This is, however, an observation about the experiments performed here and not a general property: the kinematic model provides an approximation of the response of a real arm, and the quality of that approximation varies with the movement.

Four limitations should be stated explicitly. The joint limits used in the experiments were derived from the range of the recorded reference and are therefore artificial: they served to verify that a prescribed bound is enforced, not to reproduce a clinical condition. The two movements also differed in how strongly they exercised the constraint: in the first, the unconstrained execution already remained within the imposed bounds, so that this movement did not provide a strong experimental stress test of the constraint. A single healthy subject took part, acting in both roles, so the results describe the behavior of the algorithm rather than its therapeutic effect. Finally, the barrier is applied with a fixed weight, so that the minimizer of the predicted human-space cost remains displaced towards the interior of the admissible set even when the reference is attainable. This displacement is accepted as part of the formulation. The experimentally measured trajectories exhibit an additional deviation from the reference due to the imperfect correspondence between the kinematic model and the real human response.

Overall, the work shows that patient-specific joint limits can be incorporated into the generation of the robot trajectory of an end-effector-based rehabilitation system without modifying its hardware and while preserving the kinematic structure of the method it extends. ROM compatibility is therefore addressed directly at the level at which the robot motion is generated.

8.2 Future Work

Several aspects remain open and suggest directions for further work.

1. **Clinically determined limits.** The procedure takes the admissible ranges as an input, and in the experiments reported here they were derived from the excursion of the recorded reference. The natural next step is to supply them from the clinical assessment of the patient, which requires no change to the formulation but places the method in the workflow for which it was designed. Measuring the bounds on the patient would also allow the behavior of the

algorithm to be examined when the limits are genuinely binding over a large portion of the movement, rather than over its terminal part only.

2. **Reducing the discrepancy between predicted and measured configuration.** The comparison reported in Chapter 7 shows that the configuration actually realized by the arm can differ appreciably from the one the kinematic model predicts, and that the magnitude of the difference depends on the movement. The model is parameterized from three linear measurements taken on the subject; since the motion capture system already reconstructs the position of the joint centers, the segment lengths could instead be derived from the distances between consecutive joint centers in the recorded data, obtaining a parameterization consistent with the same source from which the reference trajectory is taken. Until that agreement improves, the formulation can enforce the prescribed limits only on the predicted human configuration, while the actual joint motion must be assessed experimentally.

3. **Evaluation on patients and on a larger population.** The experiments involved a single healthy subject. A study on subjects with restricted joint ROM, and in particular on patients following a neurological event, would allow the clinical relevance of the approach to be assessed and would exercise the method in the condition it was conceived for, in which the reference systematically exceeds the admissible set.

4. **Admissible regions beyond the box.** The limits were represented as independent bounds on each joint, consistently with the common clinical practice of reporting ROM separately for individual joint motions. This remains an approximation of the actual admissible region, since the excursion available at one joint may depend on the configuration of the others. A coupled description, for example one learned from motion data [20, 32], could be introduced through general constraint functions in the barrier formulation. This extension would require the corresponding gradient and Hessian to be reformulated and would generally remove the diagonal structure exploited in Chapter 5, while preserving the overall strategy of optimizing directly in the human joint space.

Bibliography

[1] Franka Robotics, "Robot and interface specifications, Franka Control Interface documentation," https://frankarobotics.github.io/docs/control_parameters.html.

[2] Xsens Technologies B.V., *MVN User Manual*, 2021, document MVNManual, Revision Z.

[3] ——, "Euler sequences in joint angles & Gimbal lock," https://base.xsens.com/s/article/Euler-sequences-in-joint-angles-Gimbal-lock, 2022, xsens Knowledge Base article, published 14 January 2022.

[4] D. Romeo, "Human joint-driven trajectory tracking for enhanced end-effector-based robotic rehabilitation," Master's thesis, Università degli Studi di Pavia, 2025.

[5] C. Alessi, D. Romeo, N. Sacchi, and A. Ferrara, "Human joint-driven trajectory tracking for enhanced end-effector-based robotic rehabilitation," in *2026 European Control Conference (ECC)*, Reykjavík, Iceland, Jul. 2026, pp. 163–168.

[6] P. Langhorne, J. Bernhardt, and G. Kwakkel, "Stroke rehabilitation," *The Lancet*, vol. 377, no. 9778, pp. 1693–1702, 2011.

[7] H. I. Krebs and B. T. Volpe, "Rehabilitation robotics," *Handbook of Clinical Neurology*, vol. 110, pp. 283–294, 2013.

[8] J. Laut, M. Porfiri, and P. Raghavan, "The present and future of robotic technology in rehabilitation," *Current Physical Medicine and Rehabilitation Reports*, vol. 4, no. 4, pp. 312–319, 2016.

[9] P. Maciejasz, J. Eschweiler, K. Gerlach-Hahn, A. Jansen-Troy, and S. Leonhardt, "A survey on robotic devices for upper limb rehabilitation," *Journal of NeuroEngineering and Rehabilitation*, vol. 11, no. 1, p. 3, 2014.

[10] H. M. Qassim and W. Z. Wan Hasan, "A review on upper limb rehabilitation robots," *Applied Sciences*, vol. 10, no. 19, p. 6976, 2020.

[11] E. Kyrkjebø, M. J. Laastad, and Ø. Stavdahl, "Feasibility of the UR5 industrial robot for robotic rehabilitation of the upper limbs after stroke," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Madrid, Spain, Oct. 2018, pp. 1–6.

[12] B. Sheng, S. Xie, L. Tang, C. Deng, and Y. Zhang, "An industrial robot-based rehabilitation system for bilateral exercises," *IEEE Access*, vol. 7, pp. 151 282–151 294, 2019.

[13] D. Onfiani, M. Caramaschi, L. Biagiotti, and F. Pini, "Optimizing design and control methods for using collaborative robots in upper limb rehabilitation," *IEEE/ASME Transactions on Mechatronics*, vol. 30, no. 6, pp. 4849–4861, Dec. 2025.

[14] N. Sacchi, E. Vacchini, and A. Ferrara, "Human-robot ergonomic handover via deep neural network based adaptive integral sliding mode control," in *2024 European Control Conference (ECC)*, Stockholm, Sweden, Jun. 2024, pp. 585–590.

[15] C. Alessi, N. Sacchi, and A. Ferrara, "A flexible clinician-friendly interface for upper limb rehabilitation with industrial robots," in *2025 International Conference on Rehabilitation Robotics (ICORR)*. Chicago, IL, USA: IEEE, 2025, pp. 749–754.

[16] C. Alessi, N. Sacchi, C. Fundarò, and A. Ferrara, "AS-ULDAR: a ROS 2-based collaborative robotic platform for upper-limb rehabilitation — systems integration and pilot usability evaluation in a clinical environment," *Frontiers in Robotics and AI*, 2026, section Biomedical Robotics, Special Issue on Robotics and Emerging Technologies in Rehabilitation: The Robot as an Assistive Tool for Both Assessment and Personalized Intervention. Accepted October 2026, to appear.

[17] Y. Chen, G. Chen, J. Ye, X. Qiu, and X. Li, "Safe and individualized motion planning for upper-limb exoskeleton robots using human demonstration and interactive learning," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, Yokohama, Japan, 2024, pp. 15 307–15 313.

[18] C. Lauretti, F. Cordella, A. L. Ciancio, E. Trigili, J. M. Catalan, F. J. Badesa, S. Crea, S. M. Pagliara, S. Sterzi, N. Vitiello, N. Garcia Aracil, and L. Zollo, "Learning by demonstration for motion planning of upper-limb exoskeletons," *Frontiers in Neurorobotics*, vol. 12, p. 5, 2018.

[19] B. Brahmi, M. Saad, M. H. Rahman, and C. Ochoa-Luna, "Cartesian trajectory tracking of a 7-DOF exoskeleton robot based on human inverse kinematics," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 49, no. 3, pp. 600–611, Mar. 2019.

[20] S. Keyvanian, M. J. Johnson, and N. Figueroa, "Learning realistic joint space boundaries for range of motion analysis of healthy and impaired human arms," in *2023 IEEE-RAS 22nd International Conference on Humanoid Robots (Humanoids)*, Austin, TX, USA, 2023, pp. 1–8.

[21] A. M. Oosterwijk, M. K. Nieuwenhuis, C. P. van der Schans, and L. J. Mouton, "Shoulder and elbow range of motion for the performance of activities of daily

living: a systematic review," *Physiotherapy Theory and Practice*, vol. 34, no. 7, pp. 505–528, 2018.

[22] A. Ferrara, C. Alessi, and N. Sacchi, "Sliding mode control and observation for upper limb robotics rehabilitation," *International Journal of Robust and Nonlinear Control*, 2026, accepted September 2026, to appear.

[23] G. Legnani and I. Fassi, *Robotica industriale: modellazione, pianificazione, controllo, programmazione, componentistica, normativa e sicurezza*. Novara, Italy: CittàStudi, 2016.

[24] B. Siciliano, L. Sciavicco, L. Villani, and G. Oriolo, *Robotics: Modelling, Planning and Control*. Springer, 2009.

[25] J. Denavit and R. S. Hartenberg, "A kinematic notation for lower-pair mechanisms based on matrices," *Journal of Applied Mechanics*, vol. 22, no. 2, pp. 215–221, 1955.

[26] K. M. Lynch and F. C. Park, *Modern Robotics: Mechanics, Planning, and Control*. Cambridge University Press, 2017.

[27] S. R. Buss, "Introduction to inverse kinematics with Jacobian transpose, pseudoinverse and damped least squares methods," Department of Mathematics, University of California, San Diego, Tech. Rep., Oct. 2009, unpublished survey.

[28] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed. Springer, 2006.

[29] Franka Emika GmbH, "Panda data sheet: Robot arm and control," https://download.franka.de/Datasheet-EN.pdf, 2020.

[30] Franka Robotics, "franka_description: FER kinematic model," https://github.com/frankarobotics/franka_description/blob/main/robots/fer/kinematics.yaml.

[31] G. Wu, F. C. T. van der Helm, H. E. J. Veeger, M. Makhsous, P. Van Roy, C. Anglin, J. Nagels, A. R. Karduna, K. McQuade, X. Wang, F. W. Werner, and B. Buchholz, "ISB recommendation on definitions of joint coordinate systems of various joints for the reporting of human joint motion — Part II: shoulder, elbow, wrist and hand," *Journal of Biomechanics*, vol. 38, no. 5, pp. 981–992, 2005.

[32] K. Abdel-Malek and J. S. Arora, *Human Motion Simulation: Predictive Dynamics*. Academic Press, 2013.

[33] J. M. Soucie, C. Wang, A. Forsyth, S. Funk, M. Denny, K. E. Roach, and D. Boone, "Range of motion measurements: reference values and a database for comparison studies," *Haemophilia*, vol. 17, no. 3, pp. 500–507, 2011.