



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Designing Personalised Rehabilitation Controllers using Offline Model-Based Optimisation

Citation for published version:

Christou, A, Gordon, D, Stouraitis, T & Vijayakumar, S 2023, Designing Personalised Rehabilitation Controllers using Offline Model-Based Optimisation. in *Proceedings of the 2022 IEEE International Conference on Robotics and Biomimetics (IEEE ROBIO 2022)*. Institute of Electrical and Electronics Engineers (IEEE), pp. 148-155, The 2022 IEEE International Conference on Robotics and Biomimetics , Xishuangbanna, China, 5/12/22. <https://doi.org/10.1109/ROBIO55434.2022.10011916>

Digital Object Identifier (DOI):

[10.1109/ROBIO55434.2022.10011916](https://doi.org/10.1109/ROBIO55434.2022.10011916)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Proceedings of the 2022 IEEE International Conference on Robotics and Biomimetics (IEEE ROBIO 2022)

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Designing Personalised Rehabilitation Controllers using Offline Model-Based Optimisation*

Andreas Christou¹, Daniel Gordon², Theodoros Stouraitis³, and Sethu Vijayakumar⁴

Abstract—The use of robotic assistance in rehabilitation is becoming more popular, yet delivering optimal assistance remains an open challenge. In order to accelerate a patient’s recovery, assistance that is personalised to the needs of the patient is required. However, controllers of rehabilitation robots have traditionally been designed and tuned heuristically, through trial and error, with one set of parameters used across several patients. In this paper, we propose an offline model-based optimisation approach, which can be used to create personalised rehabilitation controllers. We formulate the process of designing and tuning a rehabilitation controller as a multi-objective optimisation problem, and we solve this problem using Bayesian optimisation. We evaluate our method with forward dynamics simulations and the results demonstrate that a set of controller parameters can be obtained that are both patient-specific and task-specific. Our approach could be used for the personalisation of controllers designed for rehabilitation, injury prevention and human augmentation.

I. INTRODUCTION

Rehabilitation robots have become more popular in the treatment of neurological disorders such as stroke and spinal cord injury, as they provide a reliable method for delivering physical therapy. They provide a systematic way of delivering assistance, body weight support, as well as a quantitative assessment of the patient’s performance [1], [2]. However, every person can be affected differently by a neurological disorder and the same intervention may not be optimal for everyone. The same equipment may be appropriate for helping different patients but depending on the severity of the injury, different levels of assistance, or different forms of assistance, or potentially resistance, may be needed. Designing and fine-tuning the controllers of these robots to the needs of each patient may help accelerate the patient’s recovery but it can be a very cumbersome and time-consuming process [3]. This applies to the rehabilitation of both the upper and lower limbs [4], [5]. Here we will be focusing on the personalisation of

*This research was supported in part by the Engineering and Physical Sciences Research Council (EPSRC, grant reference EP/L016834/1) as part of the Centre for Doctoral Training in Robotics and Autonomous Systems at Heriot-Watt University and The University of Edinburgh, in part by the Alan Turing Institute, U.K., and in part by the Honda Research Institute Europe, Germany.

¹Andreas Christou is with the School of Informatics, University of Edinburgh, UK, andreas.christou@ed.ac.uk

²Daniel Gordon is with the School of Informatics, University of Edinburgh, UK, daniel.gordon@ed.ac.uk

³Theodoros Stouraitis is with the Honda Research Institute Europe, Germany, and the School of Informatics, University of Edinburgh, UK, theodoros.stouraitis@honda-ri.de

⁴Sethu Vijayakumar is with the School of Informatics, University of Edinburgh, UK, and with the Alan Turing Institute, UK, sethu.vijayakumar@ed.ac.uk

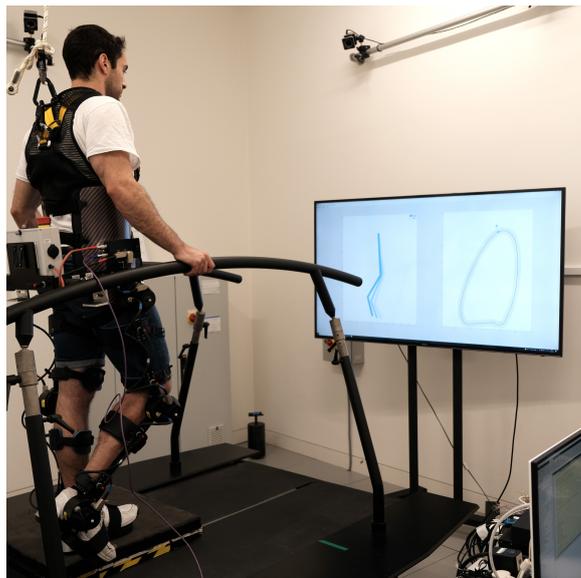


Fig. 1: Data collection setup: A healthy subject performing a trajectory tracking task while wearing a lower-limb exoskeleton.

robotic controllers for the rehabilitation of the lower limbs (Fig. 1).

In rehabilitation for locomotion, the main effort has been to help the patients learn how to reproduce a gait pattern that is healthy for them. Some of the controllers used in the early rehabilitation stage include position controllers, where the robot would strictly guide the patient along a predefined kinematic trajectory [6]. This robot-driven approach can be useful to reduce spasticity and potentially induce brain plasticity through increased somatosensory stimulation, but may also induce patient slacking and passive participation, which may prohibit or delay the recovery of patients whose residual strength is not utilised [7]. The concept of patient-driven rehabilitation was then given more attention and more controllers were studied that aimed to capture the patient’s attention and encourage them to actively participate during their therapy [6], [8]. Impedance controllers have been used for this purpose to encourage patients to use their residual strength while the robot would provide assistance only ‘as needed’. Traditionally, these controllers have been tuned heuristically, through trial and error, with one set of gains to accommodate all patients [3], [5]. While this method was found in some cases to increase the patient’s participation [9], this manual tuning of impedance controllers may result in controllers that are too stiff or too compliant for different patients, depending on the severity of their injury.

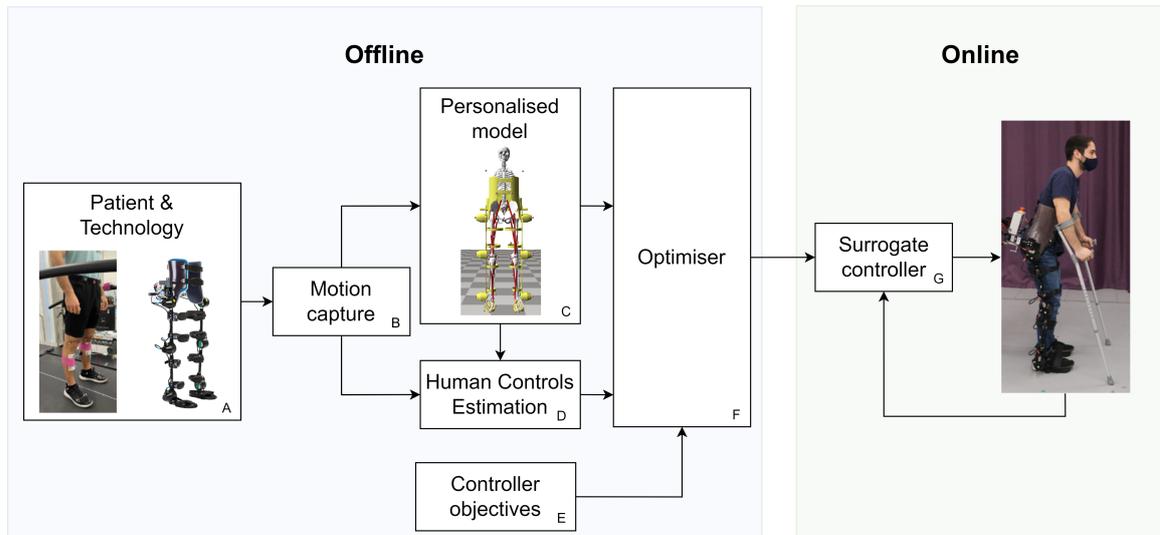


Fig. 2: Offline model-based optimisation pipeline.

Lately, the neurological rehabilitation community turned more towards methods that can provide more personalised assistance [10]–[12]. To achieve a higher level of personalisation in rehabilitation, one approach that has been proposed is the use of iterative learning controllers [13]–[15]. These controllers are based on adaptation laws which iteratively adjust one of the controller’s open parameters based on the user’s performance and the computed error. This provides a way of adjusting some of the controller’s parameters online, however it introduces additional open parameters to the controller design that need to be tuned, which requires additional testing. Another approach that has been studied in order to obtain more personalised controller parameters is human-in-the-loop optimisation [16]–[18]. In this case, a set of the controller’s open parameters are selected using Bayesian optimisation or genetic algorithms such as CMA-ES. However, this method requires the subject to perform the task multiple times while wearing the robot so that the selected algorithm can search for the optimal parameters. This can be a time-consuming process which may not be suitable, comfortable, or even safe for neurological patients.

Here, we propose a framework that uses model-based optimisation to produce personalised rehabilitation controllers, without the need for time-intensive manual tuning or excessive collection of participant data (Fig. 2). In this paper, we describe how a personalised model of the patient can be constructed along with the rehabilitation technology to be used, and how this model can be used to develop controllers based on the needs of the patient. This approach can be used for both choosing the optimal controller structure and for fine-tuning the controller’s open parameters. Here, we assume that the controller structure is known and we describe how this approach can be used to optimise the controller’s open parameters using Bayesian optimisation (BO). The optimal controller parameters obtained from this approach are then compared to a baseline impedance controller in simulation. The ability of each controller to provide assistance only

‘as needed’ for the task of following a reference kinematic trajectory is analysed.

In this paper, we make the following contributions:

- We formulate the process of designing and tuning rehabilitation controllers as an optimisation problem.
- We propose a model-based optimisation pipeline, which can be used to solve this optimisation problem offline in order to create personalised surrogate rehabilitation controllers.

In Section II a short description of the problem statement is provided. Section III includes the details of how a personalised human-exoskeleton model can be constructed and how it can be used to adjust these rehabilitation controllers to the needs of the patient. In Section IV the data collection process is described as well as the validation process of the proposed model-based optimisation approach through a series of simulation experiments. The simulation results that demonstrate the ability of the proposed method to adjust the controller parameters to the needs of the patient are presented in Section V. Section VI highlights some of the strengths and limitations of this method and Section VII provides a short summary of the contributions presented here and the way they can be used in future studies.

II. PROBLEM STATEMENT

The problem addressed here is how a personalised surrogate rehabilitation controller, H_r , can be designed based on the needs of the patient to provide personalised assistance. Given a model of the patient, P , which can include details of their musculoskeletal structure, expected behaviour, and neurological disorder, and a model of the rehabilitation robot, R , the aim is to obtain the optimal behaviour of the robot, \mathbf{u}_R , that will minimise a set of rehabilitation costs, C . This

can be expressed as:

$$H_r = \phi(\mathbf{u}_R^*), \quad (1)$$

$$\mathbf{u}_R^* = \underset{\mathbf{u}_R}{\operatorname{argmin}} f(P, R, C, \mathbf{u}_R), \quad (2)$$

where $\phi(\cdot)$ and $f(\cdot)$ are generalised functions.

The proposed model-based approach allows for the design of rehabilitation controllers to be addressed as a multi-objective optimisation problem. In this problem, the objective function contains the sum of the rehabilitation objectives, C , of the desired intervention. The equality, $g(\cdot) = 0$, and inequality constraints, $h(\cdot) > 0$, are adjusted to specify the model's dynamics as well as any restrictions associated with the controller's structure. The decision variable, \mathbf{v} , represents the controller's open parameters, that are adjusted to ensure that personalised assistance is provided. This can be expressed as:

$$\min_{\mathbf{v}} \sum_{i=1}^N w_i C_i \quad (3)$$

$$g(\mathbf{x}, \mathbf{u}) = 0, \quad (4)$$

$$h(\mathbf{x}, \mathbf{u}) > 0, \quad (5)$$

$$\mathbf{x}^- < \mathbf{x} < \mathbf{x}^+, \quad (6)$$

$$\mathbf{u}^- < \mathbf{u} < \mathbf{u}^+, \quad (7)$$

where w represents the weight of the associated cost and \mathbf{u} is the vector of controls including the estimated actions of the human and the actions of the intervention to be used. \mathbf{x} is the vector of states of the personalised model, including the estimated state of the human model and the intervention to be used, and $[\mathbf{x}^-, \mathbf{x}^+]$ and $[\mathbf{u}^-, \mathbf{u}^+]$ represent the lower and upper bounds of the model's states and controls, respectively.

The following section provides a solution to this problem. The formulation of the objective function based on the rehabilitation goals is described as well as the selection of constraints based on the constructed model and the chosen surrogate controller.

III. METHODS

A. Personalised human-exoskeleton modelling

To create a personalised human-exoskeleton model, the musculoskeletal modelling software OpenSim is used [19], [20], since it offers detailed models that can be adjusted to reflect the physical properties of the patient's body including any characteristics that may be a result of an injury or a neurological disorder. Using reflective markers and a Vicon motion capture system, a static pose of the patient is recorded (Fig. 2B), and is used to scale a generic musculoskeletal model to the patient's true size. This includes the patient's measured height and weight. The scaled model of the patient can then be combined with the right exoskeleton model to create a personalised human-exoskeleton model (Fig. 2C). The two models are combined using bushing forces at the locations of the exoskeleton's supporting cuffs to model the interaction between the cuffs and the human skin, along with any related transmission losses. The exoskeleton's limb

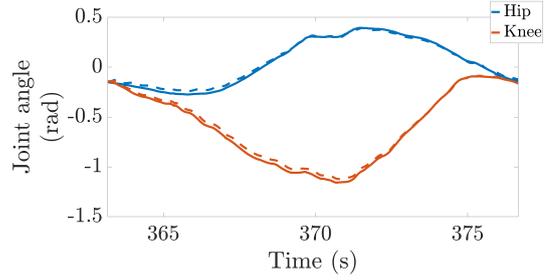


Fig. 3: Comparison between a sample recorded motion (solid line) and the motion generated through forward dynamics (dashed line).

lengths are also adjusted to ensure that the joint centre of the exoskeleton actuators aligns with the joint centre of the human joints. The coupled dynamics of this model can be expressed as:

$$\mathbf{M}_{he}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}_{he}(\mathbf{q}, \dot{\mathbf{q}}) + \mathbf{G}_{he}(\mathbf{q}) = \boldsymbol{\tau}_h + \boldsymbol{\tau}_e + \mathbf{J}(\mathbf{q})^T \mathbf{f}_{ext}, \quad (8)$$

where \mathbf{q} , $\dot{\mathbf{q}}$, $\ddot{\mathbf{q}} \in \mathbb{R}^n$ are the generalised joint positions, velocities and accelerations of the model, respectively. $\mathbf{M}_{he}(\mathbf{q}) \in \mathbb{R}^{n \times n}$ is the mass matrix of the human-exoskeleton model, $\mathbf{C}_{he}(\mathbf{q}, \dot{\mathbf{q}}) \in \mathbb{R}^n$ is the vector of Coriolis and centrifugal forces, and $\mathbf{G}_{he}(\mathbf{q}) \in \mathbb{R}^n$ is the vector of gravitational forces for a system with n degrees of freedom. $\boldsymbol{\tau}_h \in \mathbb{R}^n$ represents the human's voluntarily generated joint torques, and $\boldsymbol{\tau}_e \in \mathbb{R}^n$ are the assistive forces provided by the exoskeleton. $\mathbf{J} \in \mathbb{R}^{3 \times n}$ is the system's Jacobian and $\mathbf{f}_{ext} \in \mathbb{R}^3$ represents any external forces that may be applied to either the exoskeleton or the human model, including forces due to the human-exoskeleton coupling and ground reaction forces.

B. Human controls estimation

Using the dynamic model of the combined human-exoskeleton system, forward dynamics simulations can be carried out to predict the kinematic trajectory of the human and the exoskeleton, given the torques exerted by both. To provide an estimate of the human's expected force output, the following process is carried out.

The user of the orthosis is asked to wear the device and perform the desired task without any assistance, $\boldsymbol{\tau}_e = 0$. While performing this task, the user's motion is captured, \mathbf{q}_m . This motion is then used in a forward dynamics analysis as a reference trajectory to be followed in order to obtain the human controls, $\boldsymbol{\tau}_h$ (Fig. 2D). This is carried out using a proportional-derivative controller at the joints of the human model which can be expressed as:

$$\boldsymbol{\tau}_h = \mathbf{K}_h(\mathbf{q}_m - \mathbf{q}) + \mathbf{B}_h(\dot{\mathbf{q}}_m - \dot{\mathbf{q}}), \quad (9)$$

where \mathbf{K}_h and \mathbf{B}_h are the proportional and derivative gains, respectively. Fig. 3 shows a comparison between a sample recorded motion (solid line), and the motion generated through forward dynamics (dashed line) based on the calculated human controls, $\boldsymbol{\tau}_h$.

Given the vector of human controls, another forward dynamics simulation is carried out where the exoskeleton is

no longer in transparent mode, but provides assistive forces, $\tau_e \neq 0$. Assuming that the human strategy will not be affected by the presence of assistive forces (i.e. the vector of human controls, τ_h , obtained from the previous simulation will not change), the kinematic trajectory of the human-exoskeleton system can be predicted by:

$$\ddot{\mathbf{q}} = \mathbf{M}_{he}^{-1}(\mathbf{q})(\tau_h + \tau_e - \mathbf{C}_{he}(\mathbf{q}, \dot{\mathbf{q}}) - \mathbf{G}_{he}(\mathbf{q})). \quad (10)$$

This can then allow for the assistive forces to be adjusted in order to achieve the desired goal.

C. Exoskeleton control & offline optimisation

For this study, the goal is to design a controller where the exoskeleton will provide assistance only ‘as needed’. To achieve this, the objective function of the formulated optimisation problem, (Eq. (3)), includes two costs: a cost related to the kinematic tracking error, ϵ , and a cost related to the exoskeleton controls, \mathbf{u}_e (Fig. 2E). This can be expressed as:

$$\min_{\mathbf{v}} \frac{w_1 \sum_{i=1}^{N-1} (\mathbf{u}_{e_i}^T \mathbf{I} \mathbf{u}_{e_i})}{J_1 N - 1} + \frac{w_2 \sum_{i=1}^N \epsilon_i^2}{J_2 N}, \quad (11)$$

where w_1 and w_2 are the weights of the two costs, \mathbf{I} is the identity matrix, N is the number of the simulation time steps and J is the normalising factor. The normalising factor has been selected based on the maximum exoskeleton assistance and the maximum expected trajectory error, such that the magnitude of the two costs is comparable.

The constraints of this problem include the dynamic model as defined by Eq. (8), the upper and lower bounds of the model’s states and controls, as described by Eq. (6) - Eq. (7), and any other constraints related to the surrogate controller to be designed (Fig. 2G).

D. Surrogate controller

For this study, the path controller proposed by Duschau Wicke et al. [9] is used as an example surrogate controller. As such, the generalised constraints defined by Eq. (4)-Eq. (5) can be specified as:

$$S^* = \operatorname{argmin} \|\mathbf{Q}_{\text{ref}}(S) - \mathbf{q}_{\text{act}}\|_2, \quad (12)$$

$$\mathbf{q}_{\text{ref}} = \mathbf{Q}_{\text{ref}}(S^*), \quad (13)$$

$$\epsilon = \max(\|\mathbf{q}_{\text{ref}} - \mathbf{q}_{\text{act}}\|_2 - r_{\text{db}}, 0), \quad (14)$$

$$\Delta \tilde{\mathbf{q}} = \mathbf{q}_{\text{ref}} - \mathbf{q}_{\text{act}}, \quad (15)$$

$$\Delta q^{(j)} = \begin{cases} 0, & |\Delta \tilde{q}^{(j)}| \leq r_{\text{db}}, \\ \Delta \tilde{q}^{(j)} - r_{\text{db}}, & \Delta \tilde{q}^{(j)} > r_{\text{db}}, \\ \Delta \tilde{q}^{(j)} + r_{\text{db}}, & \Delta \tilde{q}^{(j)} < -r_{\text{db}}, \end{cases} \quad (16)$$

where $\{S^* \in \mathbb{R} : 0 < S < 1\}$ is the relative position in the gait cycle. This is defined as the position in the gait cycle where the Euclidean distance between the reference path, $\mathbf{Q}_{\text{ref}} \in \mathbb{R}^{i \times j}$, and the model’s joint configuration, $\mathbf{q}_{\text{act}} \in \mathbb{R}^j$, is minimum, where i represents the number of points in the discretised domain of the reference path and j represents the number of degrees of freedom defined by the reference path. Based on this value, the reference point, $\mathbf{q}_{\text{ref}} \in \mathbb{R}^j$, can be obtained as shown in Fig. 4. A dead band with radius,

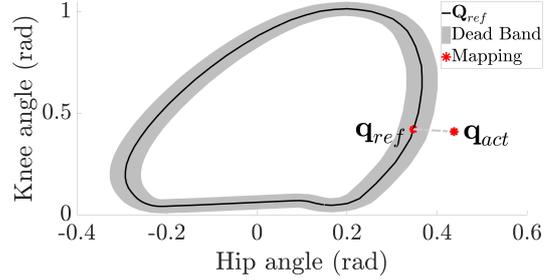


Fig. 4: Illustration of a reference kinematic path, \mathbf{Q}_{ref} , surrounded by a dead band and the mapping of the kinematic configuration of the model, \mathbf{q}_{act} , to the reference point, \mathbf{q}_{ref} , on the reference path.

r_{db} , is also defined around the reference path and is used to calculate the joint error, $\Delta \mathbf{q}$, as well as the magnitude of the error vector, ϵ .

An optional constraint that is included involves the way the relative position in the gait cycle, S , is restricted to only a constantly moving portion of the reference path, known as the *moving window*. This is to ensure that assistive forces tangential to the reference path are also provided if needed. This is achieved by defining a compliant ‘front wall’ and ‘back wall’, which are separated by a defined distance, d_w , and are constantly moving along the reference path at a speed, v_w . This speed can be obtained from the recorded motion of the user. This constraint can be expressed as:

$$v_w t - \frac{d_w}{2} \leq S \leq v_w t + \frac{d_w}{2}, \quad (17)$$

where t is the time.

Finally, we constrain the assistive controller to behave as an impedance controller. As such, the magnitude of the assistive forces can be expressed as a function of the kinematic error and the rate of change of this error. This can be expressed as:

$$\tau_e = \mathbf{K} \Delta \mathbf{q} + \mathbf{B} \Delta \dot{\mathbf{q}}, \quad (18)$$

$$\mathbf{B} = \mathbf{c}_{cr} \sqrt{\mathbf{K}}, \quad (19)$$

where \mathbf{K} and \mathbf{B} are the joint stiffness and damping of the exoskeleton’s joints, and \mathbf{c}_{cr} is the matrix of the critical damping coefficients.

E. Optimiser

Here, we describe how this problem can be solved using Bayesian Optimisation, which is a sample-efficient and gradient-free optimisation method (Fig. 2F). Using MATLAB and the *bayesopt* function, the user’s kinematic trajectory, \mathbf{q} , is predicted for various levels of exoskeleton assistance, τ_e , by adjusting the exoskeleton stiffness, \mathbf{K} . This happens in an iterative procedure where the value of the objective function (Eq. (11)) is evaluated for different exoskeleton stiffnesses to construct a model of the function (Fig. 5) and an acquisition function, which are used to choose the value of the exoskeleton stiffness to be evaluated next.

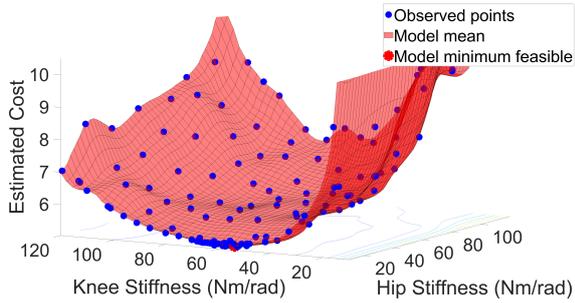


Fig. 5: The estimated objective function value as a function of exoskeleton hip stiffness and knee stiffness using Bayesian optimisation.

After N iterations, the exoskeleton stiffness, \mathbf{K}^* , that results in the lowest value of the objective function, is the output of this process along with the value of the objective function.

IV. DATA COLLECTION

A. Hardware

For this study, the exoskeleton Exo-H3 (Technaid, Spain) is used [21]. The Exo-H3 exoskeleton is a lower-limb exoskeleton with 6 active degrees of freedom, that can provide support in hip flexion and extension, knee flexion and extension, and ankle plantarflexion and dorsiflexion. Each joint is equipped with a position and velocity encoder, as well as with a torque sensor.

B. Setup

A healthy subject is asked to wear the device and perform a simplified tracking task without any assistance, $\tau_e = 0$. The user is provided with visual feedback and is instructed to follow a reference path, which is constructed based on the recorded gait of a healthy subject. The user is instructed to do so with only one leg, while the other leg supports their weight, $\mathbf{f}_{\text{ext}} = 0$ (Fig. 1). This task allows for the dimensionality of the problem and the computational demands of the optimisation to be reduced. It also eliminates the need of an additional cost to be included in the objective function, such as a cost to satisfy balance during walking. Fig. 6 shows one cycle of the trajectory obtained from a healthy subject (dashed line) trying to track the reference kinematic path (solid line).

C. Simulation experiments

To demonstrate the ability of the proposed method to find the optimal controller parameters that reflect the user's model and the user's actions, three simulation experiments were carried out: (1) a series of simulations where variation in the strength of the model is considered, (2) a series of simulations where variation in the model mass is considered, and (3) a series of simulations where variation in the model's actions is considered. Lastly, the effect of adjusting the relative weights of the objective function costs is presented in order to highlight the ability of the proposed method to provide personalised controller parameters that reflect the

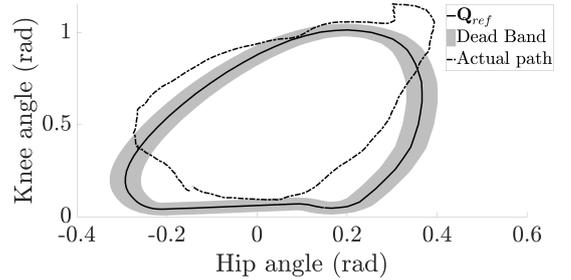


Fig. 6: A comparison between a recorded cycle (dashed line) and the reference path (solid line).

relative importance of the different rehabilitation tasks. The output of these simulations is the value of the personalised exoskeleton stiffness for the surrogate path controller (Section III-D). The personalised surrogate controller is then used to assist a simulated human and three metrics are used to evaluate the ability of the surrogate controller to provide assistance 'as needed': the normalised mean squared value of the trajectory error, hereafter referred to as *error metric*, λ_e , the normalised mean squared value of the exoskeleton assistance, hereafter referred to as *assistance metric*, λ_a , and their weighted sum, as described by Eq. (11), which is referred to as the *total rehabilitation metric*, λ_T . These values are compared to the corresponding values obtained when a baseline stiffness of magnitude $\mathbf{K} = [340, 340]$ Nm/rad [22] is used.

V. SIMULATION RESULTS

A. Personalisation

Here we demonstrate how variations in the user's musculoskeletal properties may reflect on the outputs of the model-based optimisation and how they compare to the outputs obtained from a baseline surrogate controller.

1) *Variations in muscle strength*: This experiment simulates the effects of muscle weakness, which may be the result of neurological injury. It is expected that muscle weakness will result in inadequate range of motion and thus higher assistive forces will be needed. Yet, the exoskeleton stiffness of an impedance controller that will keep the objective function value to a minimum is unknown. To investigate this, muscle weakness was simulated as a decrease in the maximum joint torque the model can generate, which results in the scaling of the input human controls, τ_h . Forward dynamics simulations were carried out for the cases where the model's maximum joint torque was scaled by 25%, 50%, 75% and 100%. Fig. 7a shows how the optimal exoskeleton stiffness and the three rehabilitation metrics change as muscle strength changes.

It can be seen that as muscle strength increases, the exoskeleton's stiffness for both the hip joint and knee joint decreases. Compared to the baseline stiffness, a higher knee stiffness and a lower hip stiffness appear to result in a lower value for the total rehabilitation metric. In all cases, it can be seen that the value of the total rehabilitation metric, λ_T , that results from the use of the personalised exoskeleton stiffness

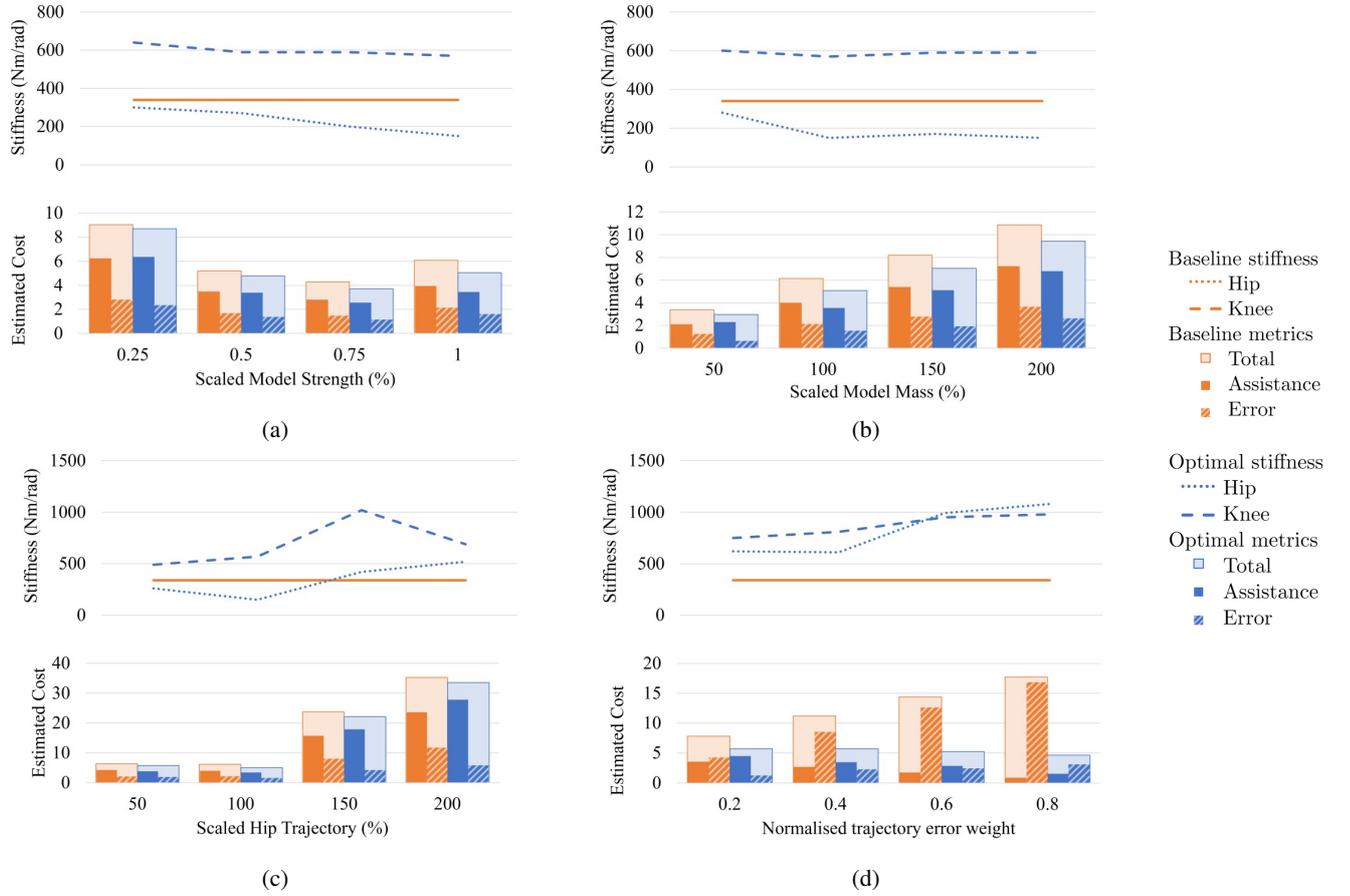


Fig. 7: Results for the simulation experiments. The dependent variables include the optimal exoskeleton stiffness, the error metric, the assistance metric and the total rehabilitation metric, while the independent variables include (a) the model’s strength, (b) the model’s mass, (c) the model’s actions and (d) the relative weights of the two costs.

is lower than the total rehabilitation metric achieved when the baseline stiffness is used.

2) *Variations in body mass:* This experiment simulates how variability in human body weight may affect the outputs of the optimiser. It is expected that as the mass of the patient increases, higher assistive forces will be needed to help them follow a reference kinematic path, however the way this reflects on the optimal exoskeleton stiffness and the value of the objective function is unclear. To test this, the mass of the personalised human-exoskeleton model was scaled to simulate individuals of lower and higher body weight. The mass of the model was scaled by 50%, 100%, 150%, and 200% and the human controls required to generate the recorded motion (Fig. 6) based on the new model mass were recomputed (Section III-B). The results from the optimisation using the updated model and human controls can be seen in Fig. 7b.

It can be seen that the value of the exoskeleton stiffness does not vary significantly with variations in the model’s mass. Yet, the value of the weighted assistance, λ_a , seems to proportionally increase as the mass of the model increases, which suggests that higher assistive forces are required to correct the motion of the model as the mass increases. It

is also evident that the value of the optimal exoskeleton stiffness is different to the baseline stiffness. An impedance controller with higher knee stiffness and lower hip stiffness than the baseline stiffness appears to result in a lower tracking error, λ_e , and a lower overall rehabilitation metric value, λ_T , for all cases.

3) *Variations in human behaviour:* This experiment simulates how variability in human controls may affect the outputs of the optimiser. Each person has a unique kinematic pattern when they walk which is reflected in the joint torques they generate. It is expected that the closest the kinematic trajectory of the patient is to the reference path, the lower the required assistance will be. However, it is not clear how the stiffness of the rehabilitation technology may need to be adjusted at the different joints given an increased trajectory error in either, or all the dimensions, of the trajectories. In these simulations, the kinematic trajectory of the recorded motion (Fig. 6), was scaled by 50%, 100%, 150%, and 200% in the direction of the hip motion and the corresponding joint torques were obtained using the human control estimation method described in Section III-B. Fig. 8 illustrates the resultant kinematic trajectory when the hip range of motion is scaled by 50% and Fig. 9 presents how the RMSE of the

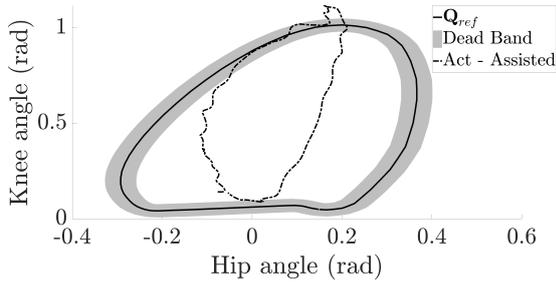


Fig. 8: The resultant kinematic trajectory that corresponds to the scaled version of the recorded motion when the range of motion of the hips is scaled by 50%.

unassisted, scaled trajectory varies over the different model actions. The results of the optimisation are presented in Fig. 7c.

From Fig. 7c it can be seen that the value of the optimal exoskeleton stiffness varies significantly over the different model actions. By comparing the variability of the optimal exoskeleton stiffness to the RMSE of the unassisted trajectories, it is evident that both the optimal exoskeleton stiffness and the assistance metric, λ_a , change according to the trajectory error generated by the different model actions. It can be seen that the RMSE of the trajectory is proportional to the optimal exoskeleton stiffness. In all cases, it is obvious that the value of both the error metric, λ_e , and the rehabilitation metric, λ_T , are lower when the optimal exoskeleton assistance is used than the values obtained for these metrics when the baseline stiffness is used.

B. Prioritising rehabilitation tasks

Depending on the severity of the injury and the patient's residual strength, different levels of assistance may be desired. By changing the weights of the costs in the objective function, the relative importance of the different rehabilitation tasks can be adjusted. In this case, these tasks include achieving accurate trajectory tracking while providing assistance only as needed. It is expected that as the relative weight of the trajectory error cost is increased, the optimal exoskeleton stiffness will also increase. Fig. 7d shows how the absolute value of the exoskeleton's stiffness, and the value of the rehabilitation metrics change when the weights of the objective function are adjusted to progressively prioritise the controller's goal of achieving accurate trajectory tracking, starting from left to right.

It can be seen that as the weight related to the trajectory tracking error increases, the optimal exoskeleton stiffness increases for both the hip joint and the knee joint. In all cases, the value of the total rehabilitation metric, λ_T , is significantly lower when the optimal exoskeleton stiffness is used than when the baseline stiffness is used. It can be noticed that in all cases, the value of the weighted assistance, λ_a , is slightly higher when the optimal stiffness is used, but this results in a significant reduction of the resultant error metric, λ_e .

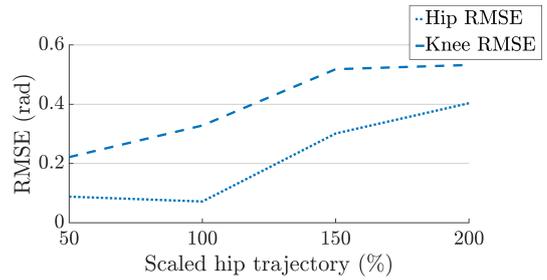


Fig. 9: RMSE of the hip and knee joint for the unassisted scaled trajectory.

VI. DISCUSSION

The presented results demonstrate that this approach enables the design of surrogate rehabilitation controllers which can be personalised to the patient's mass, residual strength and motion patterns. It is shown that for patients who are more severely affected and suffer from muscle weakness, an impedance controller with higher stiffness may be better at providing assistance 'as needed' (Fig. 7a). On the other hand, the patient's weight does not appear to have a significant impact on the optimal controller stiffness. Yet, it is evident that choosing the correct value for the controller's stiffness, based on the patient's strength and their recorded actions, is important in order to help patients with different body mass to follow the reference trajectory more accurately with less assistance (Fig. 7b). Moreover, depending on the patient's actions, a different controller stiffness may be required. It is shown that for patients whose range of motion is limited for one or more of their joints, a slightly higher controller stiffness may be required. Similarly, a higher joint stiffness may be more appropriate for patients who may produce an exaggerated motion in any direction (Fig. 7c), potentially due to spasticity or any compensatory mechanisms they may have developed. Lastly, the results indicate that based on the therapist's judgement, and which rehabilitation tasks need to be prioritised, adjustments may need to be carried out to the controller parameters. It is shown that if the ability to track a reference trajectory is considered more important than providing assistance as needed, then a higher controller stiffness may be required (Fig. 7d). Even though this result is anticipated, it is evident that the tuning of the controller's stiffness needs to be carried out for each patient independently in order to maximise the rehabilitation objectives.

The accuracy of the proposed approach is dependent on the quality of the models used, which can be considered a limitation. Accurately modelling an integrated human-exoskeleton system can be a challenge. Here, we considered the scaling of a generalised musculoskeletal model to match the user's weight and height, and combined this model with an exoskeleton model, which generates torques that perfectly match the input torque commands. However, there are several modelling details which can be further improved to more closely reflect reality, such as the maximum isometric force

of different muscle groups, the exoskeleton motor dynamics, any interaction torque losses between the human and the exoskeleton and more. Yet, given an accurate enough model, this approach provides a method for testing and fine-tuning rehabilitation controllers without the need of keeping a human in the loop.

VII. CONCLUSIONS

In this paper we presented a method where a rehabilitation path controller can be designed and fine-tuned in a simulation environment using Bayesian optimisation. The presented simulation results indicate that the same rehabilitation controller parameters may not be optimal for everyone and that with our proposed optimisation pipeline, rehabilitation controllers can be designed and fine-tuned based on the needs of the patient. We demonstrated that a different exoskeleton stiffness may be needed in order to maximise the desired rehabilitation objectives depending on the user's residual strength, the user's weight and the user's actions.

In the future, this method can be verified through experiments with neurological patients and can be used to prevent cumbersome testing with the human in the loop, where data collection and personalisation can be challenging, especially when it involves neurological patients. By adjusting (1) the input model, of either the human or the rehabilitation robot, (2) the objective function, i.e. the preferred intervention, and/or (3) the constraints, i.e. the surrogate controller's structure, this pipeline can be used to design and optimise more complicated controllers for rehabilitation, injury prevention and human augmentation. This method can also be used to study how time-dependent variables may evolve over several gait cycles and can help personalise controllers that may include variable impedance or hybrid interventions combining robotic assistance and electrical stimulation. Reducing the computation demands of the presented method may also allow the optimisation process to be carried out online and eliminate the need of creating surrogate controllers.

REFERENCES

- [1] L. Li, S. Tyson, and A. Weightman, "Professionals' Views and Experiences of Using Rehabilitation Robotics With Stroke Survivors: A Mixed Methods Survey," *Frontiers in Medical Technology*, vol. 3, no. November, 2021.
- [2] B. Hobbs and P. Artemiadis, "A Review of Robot-Assisted Lower-Limb Stroke Therapy: Unexplored Paths and Future Directions in Gait Rehabilitation," *Frontiers in NeuroRobotics*, vol. 14, no. April, 2020.
- [3] M. R. Tucker, J. Olivier, A. Pagel, H. Bleuler, M. Bouri, O. Lamberg, J. R. Del Millán, R. Riener, H. Vallery, and R. Gassert, "Control strategies for active lower extremity prosthetics and orthotics: A review," *Journal of NeuroEngineering and Rehabilitation*, vol. 12, no. 1, 2015.
- [4] S. Dalla Gasperina, L. Roveda, A. Pedrocchi, F. Braghin, and M. Gandolla, "Review on Patient-Cooperative Control Strategies for Upper-Limb Rehabilitation Exoskeletons," *Frontiers in Robotics and AI*, vol. 8, no. December, pp. 1–24, 2021.
- [5] R. Baud, A. R. Manzoori, A. Ijspeert, and M. Bouri, "Review of control strategies for lower - limb exoskeletons to assist gait," *Journal of NeuroEngineering and Rehabilitation*, pp. 1–34, 2021.
- [6] W. Meng, Q. Liu, Z. Zhou, Q. Ai, B. Sheng, and S. S. Xie, "Recent development of mechanisms and control strategies for robot-assisted lower limb rehabilitation," *Mechatronics*, vol. 31, pp. 132–145, 2015.
- [7] L. Marchal-Crespo and D. J. Reinkensmeyer, "Review of control strategies for robotic movement training after neurologic injury," *Journal of NeuroEngineering and Rehabilitation*, vol. 6, no. 1, 2009.
- [8] S. Paolucci, A. Di Vita, R. Massicci, M. Traballese, I. Bureca, A. Matano, M. Iosa, and C. Guariglia, "Impact of participation on rehabilitation results: a multivariate study," *European journal of physical and rehabilitation medicine*, vol. 48, no. 3, pp. 455–66, sep 2012.
- [9] A. Duschau-Wicke, J. Von Zitzewitz, A. Caprez, L. Lünenburger, and R. Riener, "Path control: A method for patient-cooperative robot-aided gait rehabilitation," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 18, no. 1, pp. 38–48, 2010.
- [10] Z. Sun, F. Li, X. Duan, L. Jin, Y. Lian, S. Liu, and K. Liu, "A novel adaptive iterative learning control approach and human-in-the-loop control pattern for lower limb rehabilitation robot in disturbances environment," *Autonomous Robots*, vol. 45, no. 4, pp. 595–610, 2021.
- [11] V. Rajasekaran, E. López-Larraz, F. Trincado-Alonso, J. Aranda, L. Montesano, A. J. Del-Ama, and J. L. Pons, "Volition-adaptive control for gait training using wearable exoskeleton: Preliminary tests with incomplete spinal cord injury individuals," *Journal of NeuroEngineering and Rehabilitation*, vol. 15, no. 1, pp. 1–16, 2018.
- [12] S. Maggioni, N. Reinert, L. Lünenburger, and A. Melendez-Calderon, "An adaptive and hybrid end-point/joint impedance controller for lower limb exoskeletons," *Frontiers Robotics AI*, vol. 5, no. OCT, 2018.
- [13] C. Bayón, S. S. Fricke, E. Rocon, H. Van Der Kooij, and E. H. Van Asseldonk, "Performance-Based Adaptive Assistance for Diverse Subtasks of Walking in a Robotic Gait Trainer: Description of a New Controller and Preliminary Results," *Proceedings of the IEEE RAS and EMBS International Conference on Biomedical Robotics and Biomechanics*, vol. 2018-Augus, no. 12850, pp. 414–419, 2018.
- [14] S. Srivastava, P. C. Kao, S. H. Kim, P. Stegall, D. Zanotto, J. S. Higginson, S. K. Agrawal, and J. P. Scholz, "Assist-as-Needed Robot-Aided Gait Training Improves Walking Function in Individuals Following Stroke," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 23, no. 6, pp. 956–963, 2015.
- [15] C. T. Freeman, E. Rogers, A.-m. Hughes, J. H. Burridge, and K. L. Meadmore, "Control in Health Care STROKE REHABILITATION," *IEEE Control Systems Magazine*, no. February, pp. 18–43, 2012.
- [16] D. F. N. Gordon, C. McGreavy, A. Christou, and S. Vijayakumar, "Human-in-the-Loop Optimization of Exoskeleton Assistance Via Online Simulation of Metabolic Cost," *IEEE Transactions on Robotics*, vol. 38, no. 3, pp. 1410–1429, jun 2022.
- [17] K. A. Witte, P. Fiers, A. L. Sheets-Singer, and S. H. Collins, "Improving the energy economy of human running with powered and unpowered ankle exoskeleton assistance," *Science Robotics*, vol. 5, no. 40, mar 2020.
- [18] J. Zhang, P. Fiers, K. A. Witte, R. W. Jackson, K. L. Poggensee, C. G. Atkeson, and S. H. Collins, "Human-in-the-loop optimization of exoskeleton assistance during walking," *Science*, vol. 356, no. 6344, pp. 1280–1284, jun 2017.
- [19] S. L. Delp, F. C. Anderson, A. S. Arnold, P. Loan, A. Habib, C. T. John, E. Guendelman, and D. G. Thelen, "OpenSim: Open-source software to create and analyze dynamic simulations of movement," *IEEE Transactions on Biomedical Engineering*, vol. 54, no. 11, pp. 1940–1950, 2007.
- [20] A. Seth, J. L. Hicks, T. K. Uchida, A. Habib, C. L. Dembia, J. J. Dunne, C. F. Ong, M. S. DeMers, A. Rajagopal, M. Millard, S. R. Hamner, E. M. Arnold, J. R. Yong, S. K. Lakshminathan, M. A. Sherman, J. P. Ku, and S. L. Delp, "OpenSim: Simulating musculoskeletal dynamics and neuromuscular control to study human and animal movement," *PLOS Computational Biology*, vol. 14, no. 7, p. e1006223, 2018.
- [21] M. Bortole, A. Venkatakrishnan, F. Zhu, J. C. Moreno, G. E. Francisco, J. L. Pons, and J. L. Contreras-Vidal, "The H2 robotic exoskeleton for gait rehabilitation after stroke: Early findings from a clinical study," *Journal of NeuroEngineering and Rehabilitation*, vol. 12, no. 1, pp. 1–14, 2015.
- [22] A. J. Del-Ama, Á. Gil-Agudo, J. L. Pons, and J. C. Moreno, "Hybrid FES-robot cooperative control of ambulatory gait rehabilitation exoskeleton," *Journal of NeuroEngineering and Rehabilitation*, vol. 11, no. 1, pp. 1–15, 2014.