



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Human-in-the-loop optimisation in robot-assisted gait training

Citation for published version:

Christou, A, Sochopoulos, A, Lister, E & Vijayakumar, S 2025, Human-in-the-loop optimisation in robot-assisted gait training. in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Proceedings of the International Conference on Intelligent Robots and Systems, Institute of Electrical and Electronics Engineers, pp. 18219-18225, 2025 IEEE/RSJ International Conference on Intelligent Robots and Systems, Hangzhou, China, 19/10/25. <https://doi.org/10.1109/IROS60139.2025.11246110>

Digital Object Identifier (DOI):

[10.1109/IROS60139.2025.11246110](https://doi.org/10.1109/IROS60139.2025.11246110)

Link:

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

Document Version:

Peer reviewed version

Published In:

2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)

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.



Human-in-the-loop optimisation in robot-assisted gait training

Andreas Christou¹, Andreas Sochopoulos¹, Elliot Lister¹ and Sethu Vijayakumar¹

Abstract—Wearable robots offer a promising solution for quantitatively monitoring gait and providing systematic, adaptive assistance to promote patient independence and improve gait. However, due to significant interpersonal and intrapersonal variability in walking patterns, it is important to design robot controllers that can adapt to the unique characteristics of each individual. This paper investigates the potential of human-in-the-loop optimisation (HILO) to deliver personalised assistance in gait training. The Covariance Matrix Adaptation Evolution Strategy (CMA-ES) was employed to continuously optimise an assist-as-needed controller of a lower-limb exoskeleton. Six healthy individuals participated over a two-day experiment. Our results suggest that while the CMA-ES appears to converge to a unique set of stiffnesses for each individual, no measurable impact on the subjects’ performance was observed during the validation trials. These findings highlight the impact of human-robot co-adaptation and human behaviour variability, whose effect may be greater than potential benefits of personalising rule-based assistive controllers. Our work contributes to understanding the limitations of current personalisation approaches in exoskeleton-assisted gait rehabilitation and identifies key challenges for effective implementation of human-in-the-loop optimisation in this domain.

I. INTRODUCTION

Wearable robotic devices hold great promise in enhancing the outcomes of physical therapy and reducing the physical strain on healthcare professionals. However, the growing use of robotic assistance introduces the risk of depersonalising rehabilitation, potentially losing the tailored, highly effective treatments typically provided by healthcare providers [1], [2]. To address this concern, there has been a shift toward the use of collaborative robots designed to cater to the specific needs of patients, offering personalised assistance.

In gait rehabilitation, several studies have emphasised the importance of providing “assistance as needed”—partial support that encourages the patient’s active participation [1], [3]–[5]. Although various control strategies exist to achieve this [6]–[10], it remains unclear whether there is a single control strategy that is superior to other strategies or that is in its form optimal for each individual. A prevailing issue with current practices is that many existing controllers are tuned based on the performance of a healthy participant, resulting in generalised solutions that often fail to meet individual

needs effectively. Given the significant variability in gait among individuals, it is crucial to explore new methods for adjusting control parameters in wearable devices to deliver personalised and adaptive assistance.

Human-in-the-loop optimisation (HILO) is a dynamic approach that integrates human feedback directly into the control system of wearable devices allowing for real-time adjustments and ensuring that the device adapts to the user’s unique biomechanics and physiological responses. This method leverages the cyclic nature of gait, where adjustments to robot controllers are made iteratively with each gait cycle or number of gait cycles (Figure 1a). Based on this continuous feedback loop, HILO tailors robotic controllers to the needs of the user providing individualised assistance.

To date, HILO has been successful in adjusting the assistance provided by wearable robots in order to reduce the metabolic cost of gait [11]–[14], increase the self-selected speed of walking [15], [16] in healthy subjects and reduce joint loading in manual material handling activities [17]. Using mostly one-DOF robots, studies have focused on the optimisation of parameterised assistive torque profiles for mainly the hip joint and the ankle joint. However, the effectiveness of HILO for the personalisation of rehabilitation controllers has not been studied. It is hypothesised here that HILO could be used for the personalisation of the open parameters of rehabilitation controllers in order to provide assistance as needed.

So far, to carry out HILO, two main algorithms have been used: the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) and Bayesian optimisation. Both methods, are sample-based derivative-free optimisation methods that search for the global optimum within a constrained space. However, their underlying assumptions differ which will likely affect the efficacy of HILO, particularly when applied to gait training where time-dependent gait variability is expected to be higher.

In Bayesian optimisation, a surrogate model of a continuous function is constructed based on the sampled observations. Using Gaussian process regression, the posterior probability distribution of the unknown function, $f(x)$, is iteratively updated, and is used to update the acquisition function, $a(x)$. The updated acquisition function is then used to compute the next best sample point, x_i , and this process repeats. After N observations, this process terminates and the point where the value of $f(x)$ is highest (or lowest) is obtained. To do this, a common acquisition function is the *expected improvement* function, which balances exploration and exploitation. Using the *expected improvement* acquisition function the next best sample point is computed as the point

*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 Japanese Science and Technology Agency through the Moonshot project, and in part by the Alan Turing Institute, UK, and by the Honda Research Institute Europe.

¹Andreas Christou, Andreas Sochopoulos, Elliot Lister and Sethu Vijayakumar are with the School of Informatics, University of Edinburgh, UK. andreas.christou@ed.ac.uk

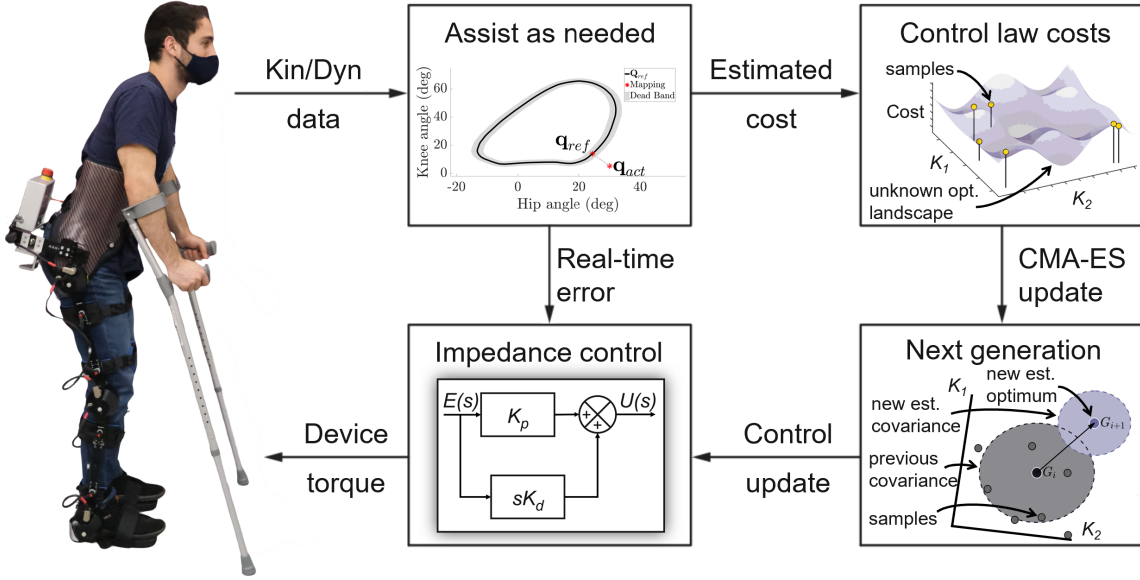


Fig. 1: HILO pipeline using the CMA-ES to personalise the open parameters of an impedance controller to provide assistance as needed.

with the highest expected quality and the highest posterior standard deviation, based on the assumption that $f(x)$ (given observations $y_{1:n}$ at points $x_{1:n}$) is normally distributed [18]. This implementation assumes noise-free evaluations and a constant function f . However, when it comes to HILO, where f represents the response of humans to an external force, some of these assumptions may be violated. It is important when using Bayesian optimisation to incorporate methods for updating the acquisition function based on realistic values for expected noise, and account for time-dependent changes in human behaviour due to fatigue, concentration and/or motor learning, which are often hard to model and predict.

In CMA-ES a stochastic search for the global optimum of an unknown function, f , is pursued iteratively through a series of generations, g . On every generation, a total of λ samples, $\{x_k^g | k \in \mathbb{N}, 1 \leq k \leq \lambda\}$, are generated by sampling a multivariate normal distribution with mean, \mathbf{m}^g , and covariance, \mathbf{C}^g , and are evaluated. Based on the observations, the new mean of the search distribution, $\mathbf{m}^{g+1} \in \mathbb{R}^n$, the new covariance matrix, $\mathbf{C}^{g+1} \in \mathbb{R}^{n \times n}$, and the new step size, $\sigma^{g+1} \in \mathbb{R}_{>0}$, are updated, where n is the dimension of the search space. A step of size, σ , in a direction dictated by the sampled observations is performed, a new sample population is generated around the new mean and this process is repeated for G generations. This process does not assume that the unknown function, f , is constant and does not prevent resampling of the same points. This allows time-dependent changes in function f to be captured and noise in the sampled observations to be accounted for.

In this work, we propose the personalisation of a lower-limb assist-as-needed controller using HILO and the CMA-ES. Through an experimental study, we observe the ability of HILO to adjust the open parameters of a gait training controller in order to help the users accurately follow a

predefined kinematic path with minimal assistance. A continuous optimisation protocol is followed over a multi-day trial as described in [13]. The results obtained from six healthy subjects are presented and discussed.

II. HILO AND ASSISTANCE AS NEEDED

Here we optimise the stiffness of an impedance controller using HILO in order to provide assistance as needed. Following the principles of path control [10], a reference kinematic path, $\mathbf{Q}_{ref} \in \mathbb{R}^{i \times 2}$, is prescribed¹ (where i is the number of points in the discretised domain of the reference path), a dead band, r_{db} , is defined around the reference path and an impedance controller is used to ensure that the user's trajectory, $\mathbf{q}_{act} \in \mathbb{R}^2$, stays close to the reference (Figure 2). Iteratively, the stiffness of the hip and the knee joints of the two legs is then adjusted and a measure of the objective function value is obtained. With the aim to provide assistance as needed, the objective function is defined as:

$$\min_{\mathbf{K}} \frac{w_1}{J_1} \frac{\sum_{i=1}^{N-1} (\mathbf{u}_{r_i}^T \mathbf{u}_{r_i})}{N-1} + \frac{w_2}{J_2} \frac{\sum_{i=1}^N (\Delta \mathbf{q}_i^T \Delta \mathbf{q}_i)}{N} + \frac{w_3}{J_3} \sum_{j=1}^4 K_j, \quad (1)$$

where the first cost aims to minimise the effort of the robot, the second aims to minimise the tracking error of the participant, and the third aims to prioritise solutions with a lower exoskeleton stiffness. The control effort cost aims to tune the assistive controller so that the exoskeleton intervenes as little as possible to the human movement, and is conflicting with the tracking cost, which aims to align human gait with the desired gait pattern. The stiffness cost

¹This reference path is used as a means of verifying the proposed HILO. It is not implied that this path is ideal for all participants.

acts as a regulariser, restricting the choice of unnecessarily large stiffness values that can lead to user discomfort.

The decision variables, $\mathbf{K} \in \mathbb{R}^4$, of the optimisation problem denote the stiffness for the hip and the knee joints, \mathbf{u}_r is the exoskeleton assistance, $\Delta \mathbf{q}$ is the kinematic tracking error, and N is the number of time steps recorded in one observation. \mathbf{w} is the vector of normalised weights for the three costs, and J is a vector of scaling factors. The scaling factors are used to normalise the cost terms to the maximum exoskeleton assistance ($\mathbf{u}_{max} = 40\text{Nm}$), the maximum expected trajectory error ($\Delta \mathbf{q}_{max} = 2^\circ$), and the maximum expected stiffness ($\mathbf{K}_{max} = 400\text{Nm/rad}$), respectively, such that the magnitude of the costs is comparable. Similarly, N is used to normalise the costs to the length of the recorded time steps and w is used to adjust the relative importance of the normalised costs. For this study, the weights in the objective function were heuristically defined to prioritise tracking performance and keep the contribution of the third cost (which favours solutions with low stiffness) to approximately 10% of the objective function value ($\mathbf{w} = [3, 1, 0.1]$).

Based on path control, the kinematic tracking error, $\Delta \mathbf{q}$, and the exoskeleton assistance, \mathbf{u}_r , are defined as:

$$\Delta \tilde{\mathbf{q}} = \mathbf{q}_{ref} - \mathbf{q}_{act}, \quad (2)$$

$$\Delta q^{(j)} = \begin{cases} 0, & |\Delta \tilde{q}^{(j)}| \leq r_{db}, \\ \Delta \tilde{q}^{(j)} - r_{db}, & \Delta \tilde{q}^{(j)} > r_{db}, \\ \Delta \tilde{q}^{(j)} + r_{db}, & \Delta \tilde{q}^{(j)} < -r_{db}, \end{cases} \quad (3)$$

$$\mathbf{u}_r = \mathbf{K} \Delta \mathbf{q} + \mathbf{B} \Delta \dot{\mathbf{q}}, \quad (4)$$

$$\mathbf{B} = \mathbf{c}_{cr} \sqrt{\mathbf{K}}, \quad (5)$$

where $\mathbf{q}_{ref} \in \mathbb{R}^2$ is the reference point that is dynamically defined as the point on the reference path that is geometrically closest to the pose of the human, and \mathbf{B} and \mathbf{c}_{cr} are the matrix of joint damping and the matrix of the critical damping coefficients of the exoskeleton, respectively.




Fig. 2: Illustration of the 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.

III. METHODOLOGY

To speed up the convergence of the optimisation through an efficient sampling method, the CMA-ES is used as described in [19]. The mean value of the search distribution

at the first generation is defined, \mathbf{m}^0 , and λ sample points, $\{x_k^1 | k \in \mathbb{N}, 1 \leq k \leq \lambda\}$ are generated based on a multivariate normal distribution with zero mean and unit variance, $\mathbf{C}^0 = \mathbf{I}$. The performance of the subjects at the generated sample points is calculated using equation 1 and the value of the mean point is updated based on a number of the sampled points, $\{\mu | \mu < \lambda\}$, which are weighted according to the subject's performance. This can be expressed as [19]:

$$\mathbf{x}_k^{g+1} \sim \mathbf{m}^g + \sigma^g \mathcal{N}(0, \mathbf{C}^g) \text{ for } k = 1, 2, \dots, \lambda, \quad (6)$$

$$\mathbf{m}^{g+1} = \mathbf{m}^g + c_m \sigma^g \sum_{i=1}^{\mu} w_i (\mathbf{x}_{i:\lambda}^{g+1} - \mathbf{m}^g), \quad (7)$$

where the symbol \sim denotes the same distribution on the left and right side, σ is the step size, c_m is the learning rate for the mean and $\mathbf{x}_{i:\lambda}^{g+1}$ is the i -th best sample out of all samples, \mathbf{x}_k^{g+1} , with the index $i : \lambda$ denoting the index of the i -th ranked sample such that $f(x_{1:\lambda}^{g+1}) \leq f(x_{2:\lambda}^{g+1}) \leq \dots \leq f(x_{\lambda:\lambda}^{g+1})$, where f is the objective function to be minimised.

On every generation, an update of the covariance matrix, \mathbf{C}^g , and the step size, σ^g , is carried out. The covariance matrix is updated such that it retains information from both the entire population, and the correlations between generations. To do this, the *cumulative* evolution path, \mathbf{p}_c^{g+1} , is utilised, which is the sequence of steps CMA-ES takes over a number of generations. This is expressed as [19]:

$$\mathbf{C}^{g+1} = (1 + c_1 \delta(h_\sigma) - c_1 - c_\mu) \mathbf{C}^g + c_1 \mathbf{p}_c^{g+1} \mathbf{p}_c^{g+1T} + c_\mu \sum_{i=1}^{\mu} w_i \mathbf{y}_{i:\lambda}^{g+1} \mathbf{y}_{i:\lambda}^{g+1T}, \quad (8)$$

$$\mathbf{p}_c^{g+1} = (1 - c_c) \mathbf{p}_c^g + h_\sigma \sqrt{c_c(2 - c_c) \mu_{eff}} \sum_{i=1}^{\mu} w_i \mathbf{y}_{i:\lambda}^{g+1}, \quad (9)$$


$$\mathbf{y}_{i:\lambda}^{g+1} = (\mathbf{x}_{i:\lambda}^{g+1} - \mathbf{m}^g) / \sigma^g, \quad (10)$$

where c_1 , c_μ and c_c are the learning rates for the rank-one and rank- μ updates of the covariance matrix and the *cumulative* evolution path, respectively, μ_{eff} is the effective sample size of the selected samples defined as $\mu_{eff} = 1 / \sum_{i=1}^{\mu} w_i^2$, $\delta(h_\sigma)$ is defined as $\delta(h_\sigma) = (1 - h_\sigma) c_c (2 - c_c) \leq 1$, and h_σ is a Heaviside function that stalls the update of the *cumulative* evolution path depending on the size of the *conjugate* evolution path. The *conjugate* evolution path, \mathbf{p}_σ^g , is independent of the direction of the successive steps performed, and is used to update the step size. The function, h_σ , and the *conjugate* evolution path are defined as [19]:

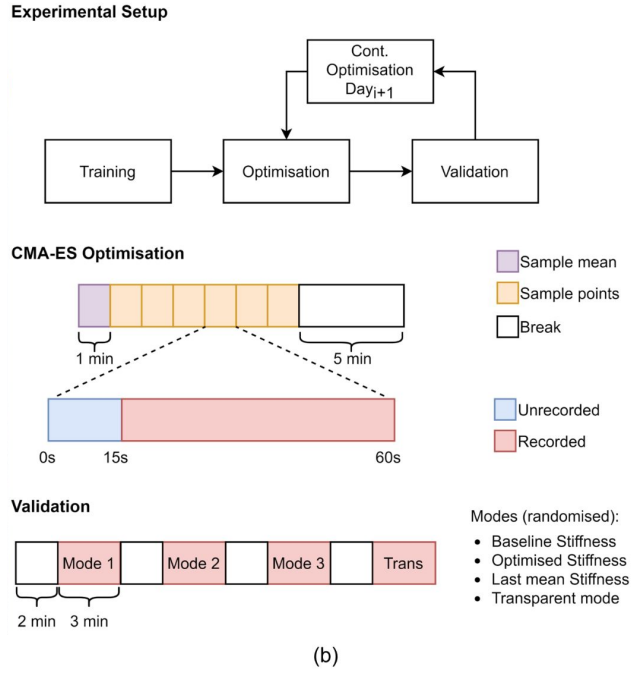
$$h_\sigma = \begin{cases} 1, & \text{if } \frac{\|\mathbf{p}_\sigma^{g+1}\|}{\sqrt{1 - (1 - c_\sigma)^{2(g+1)}}} < (1.4 + \frac{2}{n+1}) \mathbb{E} \|\mathcal{N}(\mathbf{0}, \mathbf{I})\|, \\ 0, & \text{otherwise,} \end{cases} \quad (11)$$

$$\mathbb{E} \|\mathcal{N}(\mathbf{0}, \mathbf{I})\| \approx \sqrt{n} (1 - \frac{1}{4n} + \frac{1}{21n^2}), \quad (12)$$

$$\mathbf{p}_\sigma^{g+1} = (1 - c_\sigma) \mathbf{p}_\sigma^g + \sqrt{c_\sigma(2 - c_\sigma) \mu_{eff}} \mathbf{C}^{g-\frac{1}{2}} \sum_{i=1}^{\mu} w_i \mathbf{y}_{i:\lambda}^{g+1}, \quad (13)$$



(a)



(b)

Fig. 3: (a) Healthy participant walking on a self-paced treadmill with real-time visual feedback and assistance from the exoskeleton, Exo-H3. (b) Experimental protocol for HILO following a continuous optimisation protocol over multiple days.

where c_σ is the learning rate for the *conjugate* evolution path.

Using the *conjugate* evolution path, the step size of the next generation is adjusted. This is to ensure a faster convergence and either increase the step size if the steps recorded are pointing in the same direction or decrease the step size if the steps recorded are not converging and move in opposite directions. This is achieved by comparing the length of the *conjugate* evolution path, \mathbf{p}_σ^{g+1} , with its expected length, $\mathbb{E}\|\mathcal{N}(\mathbf{0}, \mathbf{I})\|$. The adaptation of the step size can be expressed as [19]:

$$\sigma^{g+1} = \sigma^g \exp\left(\frac{c_\sigma}{d_\sigma} \left(\frac{\|\mathbf{p}_\sigma^{g+1}\|}{\mathbb{E}\|\mathcal{N}(\mathbf{0}, \mathbf{I})\|} - 1\right)\right), \quad (14)$$

where d_σ is a damping parameter.

The implementation of CMA-ES used for this study is summarised in Algorithm 1.

IV. EXPERIMENTAL VALIDATION

A. Subjects

The effectiveness of the HILO was tested on six healthy subjects (age = 30 ± 4.6 , weight = 73.4 ± 10.5 kg, 2 females). All participants were first-time users of a wearable robot. The experiment pipeline was approved by the University of Edinburgh, School of Informatics Ethics Committee (ID 2021/46920) and the participants provided written consent.

B. Hardware

The instrumented treadmill M-Gait (Motek Medical, Netherlands) was used to enable self-paced gait during the experiment and the exoskeleton Exo-H3 (Technaid, Spain) was used to provide assistance during gait (an upgrade of the

Algorithm 1: Pseudocode for CMA Evolution Strategy

Input: $0 < w < 1$

$\mathbf{C} \leftarrow \mathbf{I}$, $\mathbf{p}_c \leftarrow \mathbf{0}$, $\mathbf{p}_\sigma \leftarrow \mathbf{0}$, $g \leftarrow 0$, $\mathbf{m} \leftarrow \frac{1}{2} \mathbf{K}_{max}$

while $g < G$ **do**

```

     $\mathbf{x}^g \leftarrow \text{sample\_population}(\mathbf{m}, \sigma, \mathbf{C})$ 
     $f(\mathbf{x}^g) \leftarrow \text{evaluate\_population}(\text{HIL experiments})$ 
     $\mathbf{x}_{i:\lambda}^g \leftarrow \text{sort\_population}(\mathbf{x}^g, f(\mathbf{x}^g))$ 
     $\mathbf{m}^{g+1} \leftarrow \text{update\_mean}(\mathbf{m}^g, w, \mathbf{x}_{i:\lambda}^g, \sigma)$ 
     $\mathbf{p}_\sigma^{g+1} \leftarrow \text{update\_conj\_path}(\mathbf{p}_\sigma^g, \mathbf{C}, \mathbf{m}^g, w, \mathbf{x}_{i:\lambda}^g, \sigma)$ 
     $\sigma^{g+1} \leftarrow \text{update\_step\_size}(\sigma^g, \mathbf{p}_\sigma^{g+1})$ 
     $\mathbf{p}_c^{g+1} \leftarrow \text{update\_cum\_path}(\mathbf{p}_c^g, \mathbf{m}^g, w, \mathbf{x}_{i:\lambda}^g, \sigma)$ 
     $\mathbf{C}^{g+1} \leftarrow \text{update\_C}(\mathbf{C}^g, \mathbf{p}_c^{g+1}, \mathbf{m}^g, w, \mathbf{x}_{i:\lambda}^g, \sigma)$ 
     $\mathbf{S} \leftarrow \{\mathbf{S}; [\mathbf{x}^g, f(\mathbf{x}^g)]\}$ 
     $g \leftarrow g + 1$ 
  
```

$\mathbf{x}^* \leftarrow \mathbf{S}$ such that $f(\mathbf{x})^* = \min \mathbf{S}$

version Exo-H2 presented in [20]). The exoskeleton's joint position sensors were used to record the joint angles of the legs and provide real-time visual feedback to the user (Fig. 3a). Simulink Desktop Real Time was used for the real-time control of the exoskeleton at 100 Hz.

C. Experimental Setup

The experiments involved assisted gait training at different levels of exoskeleton stiffness. The participants were fitted with the exoskeleton and were asked to walk on the treadmill at their preferred speed in order to track the reference path as accurately as possible. The recorded kinematics of a healthy

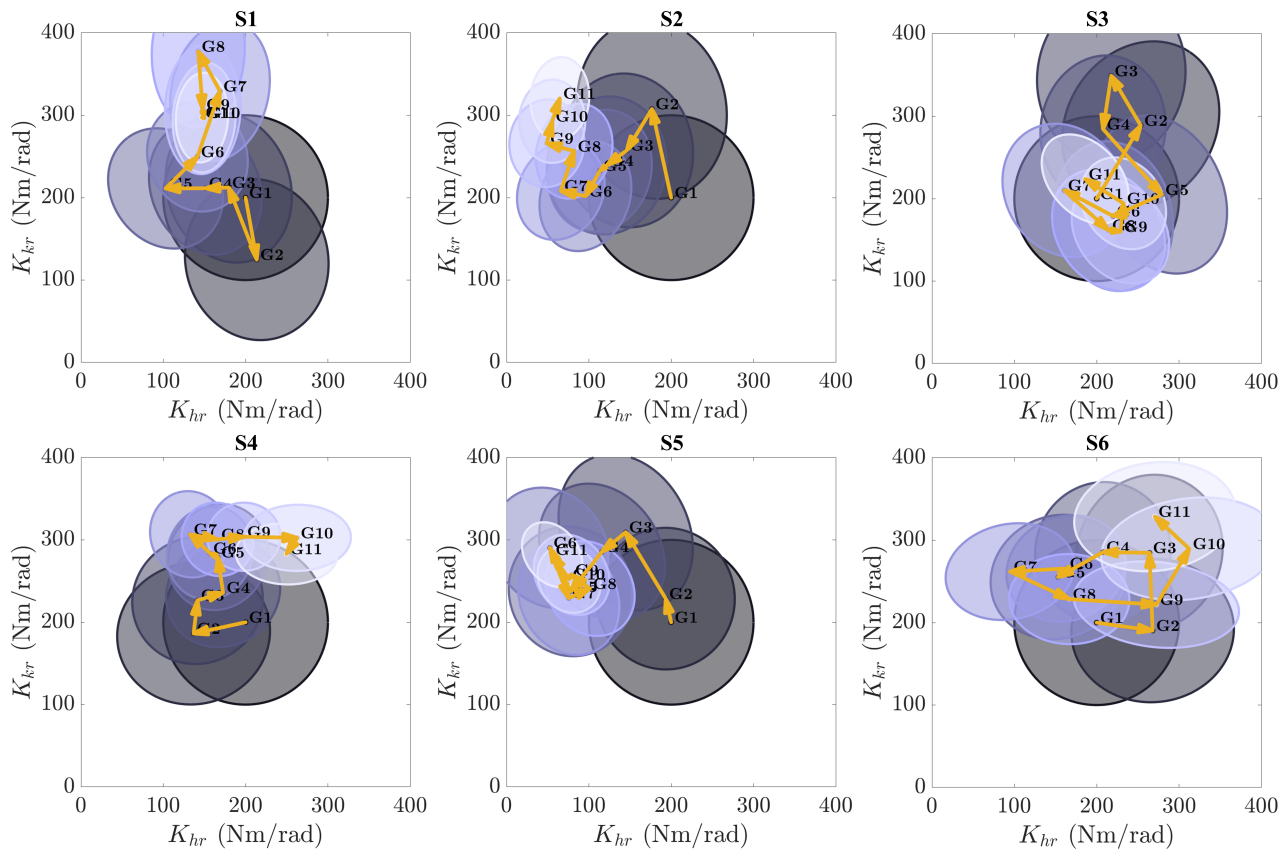


Fig. 4: Adaptation of covariance matrix and the CMA-ES generation mean, indicate from G1-G11, for all generations and all subjects. Yellow arrows show the CMA-ES step size and direction, and map out the progression of the generation mean.

subject were used as the reference path, and the path was adjusted to a path of a less pronounce loading response. Adjustments to the reference path were also made to increase comfort for each subject.

Prior to optimisation, a training period was included to familiarise the subjects with the task and the visual feedback. At all times, real-time visual feedback was provided to allow the participants to compare their kinematics to the reference, and inform them about the remaining duration of the experiment and their performance. During optimisation, a new exoskeleton stiffness was tested every minute and the performance of the subjects was measured as described by equation 1. To reduce bias from the changing exoskeleton stiffness, the first 15 seconds of each trial were discarded. One generation of the CMA-ES included 7 sample points, making up a 7-minute bout. After each 7-minute bout, a 5-minute break was allowed to reduce bias from fatigue. A total of 5 generations were performed per day, resulting in a total of 10 generations, or 70 sampled points, over the two-day experiment. At the end of each day three rounds of 3-minute validation trials were carried out in a randomised order to evaluate the effectiveness of the HILO. These included a trial with no assistance, a trial with the baseline stiffness (defined as $K_0 = 200\text{Nm/rad}$), a trial with the best identified stiffness, and a trial with the last mean stiffness obtained from the

last CMA-ES generation (Figure 3b). Lastly, a trial with no assistance was carried out.

To reduce the risk of overwhelming the participants with sensory information, subjects were asked to carry out the experiment on only one of the two legs (while the other leg was controlled with a constant low stiffness $K = 50\text{ Nm/rad}$). The convergence of the HILO and its effect on the performance of the participants are observed and discussed.

D. Analysis

Statistical analysis was carried out using two-way repeated-measures ANOVA with controller type (baseline, optimised, and last mean) and time (day 1 and day 2) as within-subject factors. Additionally, a one-way repeated measures ANOVA was performed to analyse differences between controller types regardless of time. Shapiro-Wilk's test of normality was applied to model residuals, and Mauchly's test was conducted to evaluate sphericity. If the sphericity assumption was violated, epsilon corrections were applied. Following the ANOVA, post-hoc pairwise comparisons were performed using MATLAB's *multcompare* function with Bonferroni correction to account for multiple comparisons.

V. RESULTS

Figure 4 shows the adaptation of the CMA-ES for all subjects. A continuous adaptation of the CMA-ES can be

TABLE I: Hyperparameter values for the controller and optimisation algorithm.

Controller Hyperparameters		
Max Exoskeleton assistance	\mathbf{u}_{max}	40Nm
Max expected stiffness	\mathbf{K}_{max}	400Nm/rad
Max trajectory error	$\Delta\mathbf{q}_{max}$	2
Baseline stiffness	\mathbf{K}_0	200Nm/rad
Critical damping	c_r	10
Effort cost weight	w_1	3
Tracking cost weight	w_2	1
Stiffness cost weight	w_3	0.1
CMA-ES Hyperparameters		
Sample points	λ	7
Step size	σ_0	150
LR - mean	c_m	1
LR - rank-one covariance	c_1	0.15
LR - rank- μ covariance	c_μ	0.058
LR - rank- σ covariance	c_σ	0.62

seen. In most cases this led to an assistive controller which included a higher stiffness at the knee joint compared to both the hip joint and the baseline stiffness. This suggests that despite the variability in gait between and within individuals, for the majority of the 2-day trial the participants were able to follow the reference path more accurately when the stiffness of the knee joint was higher than the baseline stiffness and the hip stiffness. It can also be seen that some participants were more consistent with their performance than others. This is revealed by the direction in which the mean points of the CMA-ES progress. For example, the progression of the generated mean points for subjects S1, S2, S4 and S5 appears to be smoother and more gradual, than subjects S3 and S6. This also led to a gradually decreasing covariance which is less evident in the results obtained for subjects S3 and S6.

In Figures 5-6 the results from the validation trials are presented. It can be seen that the performance of the subjects has significantly improved from day 1 to day 2 with all controllers (Figure 5). However, no significant differences were observed between groups when the performance of the subjects was compared when using the baseline stiffness, the optimised stiffness or the stiffness at the last mean point of the CMA-ES (Figure 6). High performance variability was observed both between and within subjects, which may have detracted from the benefits due to the adjusted controller stiffness.

VI. DISCUSSION

Here, an implementation of HILO is proposed for the personalisation of robot-assisted gait training. For this purpose the CMA-ES is used and a continuous optimisation experimental protocol is suggested to allow enough time for the optimiser to converge over multiple days. This aligns with the needs of gait training, where gait trials may need to be interrupted for rest and safety. To the best of our knowledge, this is the first time HILO has been tested for the personalisation of assistive controllers in gait rehabilitation and for robots with multiple DOFs. It is hypothesised that

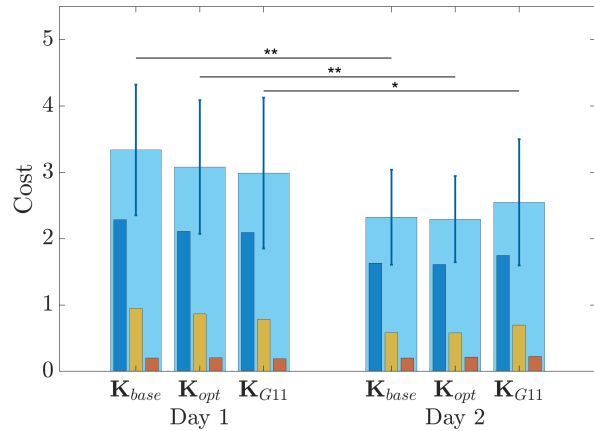


Fig. 5: Results from validation trials with the baseline stiffness, \mathbf{K}_{base} , the optimised stiffness, \mathbf{K}_{opt} , and the last mean stiffness, \mathbf{K}_{G11} , for day 1 and day 2 separately. Error bars show the standard deviation of cost (* $P < 0.05$, ** $P < 0.01$).

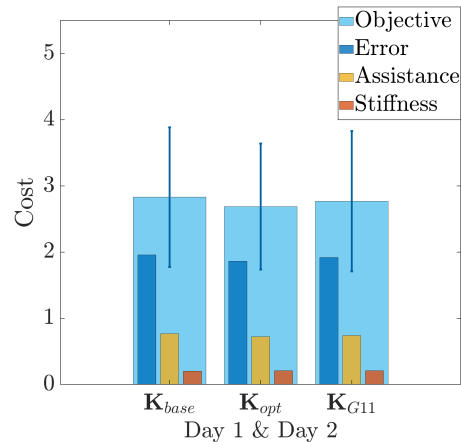


Fig. 6: Results from validation trials with the baseline stiffness, \mathbf{K}_{base} , the optimised stiffness, \mathbf{K}_{opt} , and the last mean stiffness, \mathbf{K}_{G11} , for day 1 and day 2 combined. Error bars show the standard deviation of cost.

with the human in the loop, the robot controller can be iteratively optimised based on the user's performance in order to provide assistance as needed. This could in turn allow ambulatory patients to benefit from robot-assisted gait training in order to increase independence and gait efficiency.

Results from six healthy subjects are obtained indicating that the CMA-ES can adapt the stiffness of the robot controller based on the performance of the subjects to personalise gait training. However, conclusive results regarding the effectiveness of HILO in providing assistance as needed could not be obtained. One major challenge in HILO, and in the personalisation of closed-loop robotic controllers, is intrapersonal variability. Variability within individuals may be both time-dependent and random, which poses a significant challenge in the design of personalised interventions. For an adaptive control algorithm to prove effective, the

time-dependent changes in human behaviour need to be captured and the benefits of adaptation need to outweigh any effects of unpredicted behaviour variability. This appears to be especially hard when the adaptation of rule-based controllers is considered, since such controllers are expected to be effective within a range of control parameters.

A key factor contributing to this variability is the dynamic interaction between human and robot, where both systems continuously adapt to each other. On the one hand, adjustments to the stiffness of the robot are carried out to provide support to the human in an optimal way, but on the other hand, the human behaviour concurrently changes to utilise this stiffness perhaps for a different subject-specific objective (e.g. comfort or metabolic efficiency). This co-adaptation, influenced by factors such as fatigue, comfort, motor learning, and concentration, introduces additional complexity and bias into the optimisation process that may outweigh the intended benefits of robot adaptation. In healthy individuals, this voluntary adaptation is likely more pronounced, as they have greater control over their lower limbs and can actively respond to changes in assistance. Conversely, neurological patients with impaired motor function may exhibit less voluntary adaptation, potentially reducing optimisation bias due to co-adaptation. However, this advantage is counterbalanced by other challenges. Neurological patients are more prone to fatigue, unintentional motor commands, and increased variability in movement patterns, all of which can introduce noise into the optimisation process and hinder convergence.

VII. CONCLUSION

Our work contributes to understanding the limitations of current personalisation approaches in robot-assisted gait rehabilitation and identifies key challenges for effective implementation of human-in-the-loop optimisation in this domain. Our results from six individuals suggest that while the CMA-ES can continuously adapt the stiffness parameters over time, the impact of these adaptations on performance may be masked by the high degree of human variability and co-adaptation. Future studies can focus on refining optimisation protocols to better account for human behavioural variability in order to improve the effectiveness of HILO. This may include more noise-resilient algorithms or hybrid approaches that combine HILO with therapist input or model-based constraints which may be able to more accurately account for the human-robot co-adaptation. Tailoring experimental protocols to patient-specific needs, whether by adjusting trial durations, refining performance metrics, or mitigating fatigue-related biases could also prove beneficial. Moreover, identifying combinations of rehabilitation tasks and assistive controllers where personalisation is likely to yield substantial gains is critical, particularly in cases where controller performance is highly sensitive to parameter selection, as opposed to strategies that are inherently more robust across a wider range of parameters. Finally, investigating HILO in clinical populations, such as individuals post-stroke where voluntary adaptation is more constrained, could offer clearer insight into its practical benefits and play a key role in realising the

full potential of personalisation to improve clinical outcomes and accelerate recovery in robot-assisted rehabilitation.

REFERENCES

- [1] A. Pennycott, D. Wyss, H. Vallery, V. Klamroth-Marganska, and R. Riener, "Towards more effective robotic gait training for stroke rehabilitation: A review," *Journal of NeuroEngineering and Rehabilitation*, vol. 9, pp. 1–13, 2012.
- [2] J. F. Israel, D. D. Campbell, J. H. Kahn, and T. G. Hornby, "Metabolic Costs and Muscle Activity Patterns During Robotic- and Therapist-Assisted Treadmill Walking in Individuals With Incomplete Spinal Cord Injury," *Physical Therapy*, vol. 86, no. 11, 2007.
- [3] S. Paolucci, A. D. Vita, R. Massicci, M. Traballasi, 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, pp. 455–66, 9 2012.
- [4] A. Kaelin-Lane, L. Sawaki, and L. G. Cohen, "Role of voluntary drive in encoding an elementary motor memory," *Journal of Neurophysiology*, vol. 93, pp. 1099–1103, 2005.
- [5] M. Lotze, C. Braun, N. Birbaumer, S. Anders, and L. G. Cohen, "Motor learning elicited by voluntary drive," *Brain*, vol. 126, pp. 866–872, 2003.
- [6] E. Shahriari, D. Zardykhan, A. Koenig, E. Jensen, and S. Haddadin, "Energy-based adaptive control and learning for patient-aware rehabilitation," *IEEE International Conference on Intelligent Robots and Systems*, pp. 5671–5678, 2019.
- [7] A. Taherifar, G. Vossoughi, and A. S. Ghafari, "Variable admittance control of the exoskeleton for gait rehabilitation based on a novel strength metric," *Robotica*, vol. 36, pp. 427–447, 2018.
- [8] S. Hussain, P. K. Jamwal, M. H. Ghayesh, and S. Q. Xie, "Assist-as-needed control of an intrinsically compliant robotic gait training orthosis," *IEEE Transactions on Industrial Electronics*, vol. 64, no. 2, 2017.
- [9] A. E. Russi and M. A. Brown, "Robotic assist-as-needed as an alternative to therapist-assisted gait rehabilitation?" *Int J Phys Med Rehabil*, vol. 165, pp. 255–269, 2016.
- [10] A. Duschau-Wicke, A. Caprez, and R. Riener, "Patient-cooperative control increases active participation of individuals with sci during robot-aided gait training," *Journal of NeuroEngineering and Rehabilitation*, vol. 7, pp. 1–13, 2010.
- [11] 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, pp. 1410–1429, 6 2022.
- [12] 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, pp. 1280–1284, 6 2017.
- [13] K. L. Poggensee and S. H. Collins, "How adaptation, training, and customization contribute to benefits from exoskeleton assistance," *Science Robotics*, vol. 6, p. 2021.04.25.440289, 9 2021.
- [14] G. M. Bryan, P. W. Franks, S. Song, A. S. Voloshina, R. Reyes, M. P. O'Donovan, K. N. Gregorczyk, and S. H. Collins, "Optimized hip-knee-ankle exoskeleton assistance at a range of walking speeds," *Journal of NeuroEngineering and Rehabilitation*, vol. 18, 12 2021.
- [15] S. Song and S. H. Collins, "Optimizing exoskeleton assistance for faster self-selected walking," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 29, pp. 786–795, 2021.
- [16] P. Slade, M. J. Kochenderfer, S. L. Delp, and S. H. Collins, "Personalizing exoskeleton assistance while walking in the real world," *Nature*, vol. 610, pp. 277–282, 2022.
- [17] A. Sochopoulos, T. Poliero, D. Caldwell, J. Ortiz, and C. Di Natali, "Human-in-the-loop optimization of active back-support exoskeleton assistance via lumbosacral joint torque estimation," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 6090–6096.
- [18] P. I. Frazier, "A tutorial on bayesian optimization," pp. 1–22, 2018.
- [19] N. Hansen, "The cma evolution strategy: A tutorial," 4 2016.
- [20] 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, pp. 1–14, 2015.