

Learning Personalised Human Sit-to-Stand Motion Strategies via Inverse Musculoskeletal Optimal Control

Daniel F. N. Gordon, Andreas Christou, Theodoros Stouraitis, Michael Gienger, and Sethu Vijayakumar

Abstract—Physically assistive robots and exoskeletons have great potential to help humans with a wide variety of collaborative tasks. However, a challenging aspect of the control of such devices is to accurately model or predict human behaviour, which can be highly individual and personalised. In this work, we implement a framework for learning subject-specific models of underlying human motion strategies using inverse musculoskeletal optimal control. We apply this framework to a specific motion task: the sit-to-stand transition. By collecting sit-to-stand data from 4 subjects with and without perturbations, we show that humans modulate their sit-to-stand strategy in the presence of instability, and learn the corresponding models of these strategies. In the future, the personalised motion strategies resulting from this framework could be used to inform the design of real-time assistance strategies for human-robot collaboration problems.

I. INTRODUCTION

Robots and exoskeletons are increasingly being employed in collaborative tasks to assist humans in the workplace [1], [2], and as medical devices for rehabilitation [3] or locomotion assistance [4]. These scenarios can be thought of as shared control problems, in which predicting the actions of the uncontrolled human agent is important to ensure the efficacy of the system as a whole [5]. Biomechanical models of the human body can provide insight as to the physiological dynamics that underpins motions [6], but do not easily account for individual biases which can lead to subject-specific motion strategies [7].

In this paper, we present and validate a framework for learning personalised human motion strategies for specific tasks using inverse musculoskeletal optimal control. As an example application of this framework, we learn strategies for sit-to-stand transitions (Figure 1) using experimental data collected from 4 subjects. This choice of motion task is motivated by the prevalence of musculoskeletal injuries affecting nurses and other healthcare professionals due to high exertion at work, and the increasing interest in occupational exoskeletons to circumvent this issue [8].

Other works have considered the use of biomechanical models to simulate human motion strategies, however typically the relative weighting of the objective function is fixed

This work was supported by the Alan Turing Institute, UK and Honda Research Institute Europe GmbH, Germany.

Daniel F. N. Gordon and Sethu Vijayakumar are with the School of Informatics, The University of Edinburgh, Edinburgh, UK, and also with the Alan Turing Institute, London, UK (e-mail: daniel.gordon@ed.ac.uk, sethu.vijayakumar@ed.ac.uk).

Andreas Christou is with the School of Informatics, The University of Edinburgh, Edinburgh, UK (e-mail: andreas.christou@ed.ac.uk).

Theodoros Stouraitis and Michael Gienger are with Honda Research Institute Europe GmbH, Offenbach am Main, Germany (e-mail: theodoros.stouraitis@honda-ri.de, michael.gienger@honda-ri.de).

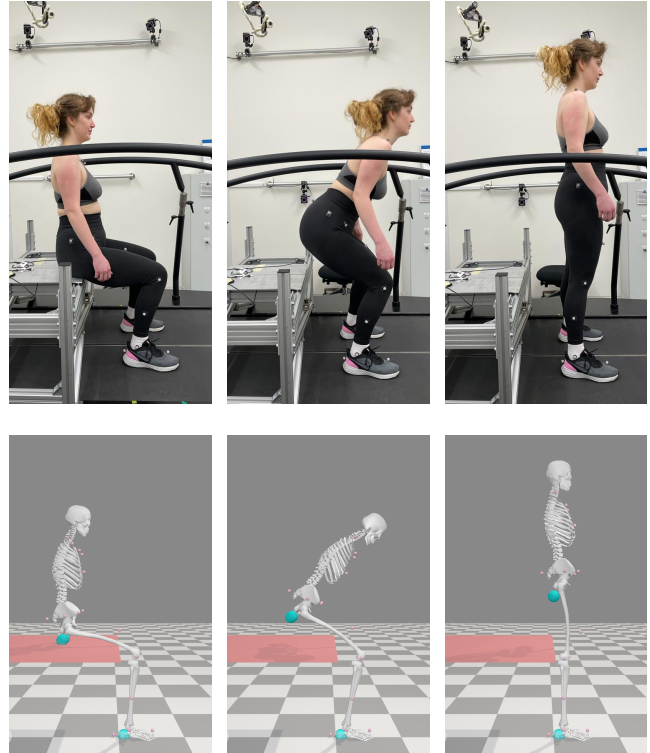


Fig. 1. **Upper row:** snapshots of a subject completing a sit-to-stand transition. **Lower row:** corresponding images of a reconstruction of the movement using a musculoskeletal model. Blue spheres represent contact geometries. The red plane represents a seat, which acts as a contact surface and can be placed at a specified height.

[9]–[11], i.e. there is no inverse step, and personalisation would require hand-tuning weights for individual subjects. Those approaches which do include an inverse step differ from our approach in the consideration of different motion tasks [12], [13] or optimisation approaches [14], [15]. In addition, our paper provides additional insights in to biomechanical adaptations during perturbed sit-to-stand transfers.

To summarise, the key contributions of this paper are the development of an inverse musculoskeletal optimal control framework for learning personalised human motion strategies, and the accompanying validation of this methodology against simulated and experimental human sit-to-stand data. In Section II we introduce and formulate the inverse musculoskeletal control framework. In Section III we instantiate a particular instance of the framework, applied to the sit-to-stand transfer motion, and discuss the specifics of the musculoskeletal model and constraints. In Section IV we present a computational validation of the framework, and in

Section V we examine the effect of perturbations on underlying human sit-to-stand motion strategies by applying the framework to experimental data collected from 4 subjects. Finally, in Section VI, we discuss key features of the results and highlight areas for further research.

II. INVERSE MUSCULOSKELETAL OPTIMAL CONTROL

Inverse optimal control is a bi-level optimisation technique used to identify cost functions that correspond to measured data samples [13]. The approach is underpinned by the assumption that the observed data is optimal with respect to some underlying objective function, and thus the technique is of particular relevance for the study of human locomotion [16]. Structurally, a lower-level optimal control optimisation enforces the dynamics and constraints, while an upper-level gradient-free optimisation modifies the relative composition of the objective to match the experimental samples.

The inverse optimal control framework presented in this work consists of a forward musculoskeletal optimal control problem in the lower-level wrapped by a sampling based upper-level optimiser. We discuss the details of each level in turn in the forthcoming subsections.

A. Musculoskeletal Optimal Control

Simulation-based investigations are increasingly employed by the biomechanics community to gain insights in to human locomotion [17]. In particular, OpenSim [6] is a widely-used open-source software package which allows for the building and analysis of custom musculoskeletal models for movement analysis [18]. In this work, we employ OpenSim Moco [19], an open-source optimal control solver which builds upon OpenSim to enable solving optimal control problems with musculoskeletal models. Assuming a known finite time horizon $[0, T]$ and constrained initial and final states \mathbf{x}_0 and \mathbf{x}_T , respectively, we consider the following musculoskeletal optimal control problem¹:

$$\begin{aligned}
& \min_{\mathbf{u}} J(\mathbf{x}(t), \mathbf{u}(t)), \\
& J(\mathbf{x}(t), \mathbf{u}(t)) = \sum_{i=1}^n w_i G_i(\mathbf{x}(t), \mathbf{u}(t)), \\
& G_i(\mathbf{x}(t), \mathbf{u}(t)) = \int_0^T g_i(\mathbf{x}(t), \mathbf{u}(t), t) dt, \\
& \mathbf{x}(t) = (\mathbf{q}(t), \dot{\mathbf{q}}(t), \mathbf{z}(t)), \\
& \dot{\mathbf{x}}(t) = \mathbf{f}_m(\mathbf{x}(t), \mathbf{u}(t), t), \\
& \theta(\mathbf{q}) = 0, \\
& \mathbf{x}(0) = \mathbf{x}_0, \\
& \mathbf{x}(T) = \mathbf{x}_T, \\
& \mathbf{x}_L \leq \mathbf{x} \leq \mathbf{x}_U, \\
& \mathbf{u}_L \leq \mathbf{u} \leq \mathbf{u}_U,
\end{aligned} \tag{1}$$

which hereafter we refer to as System 1. Here, the state vector \mathbf{x} is composed of generalised coordinates \mathbf{q} , generalised

velocities $\dot{\mathbf{q}}$, and additional state variables \mathbf{z} induced by the use of muscle-based actuators (e.g. muscle fiber lengths), while \mathbf{u} are the system controls. The function \mathbf{f}_m represents musculoskeletal dynamics, and we have bounds on both states and controls. The function θ imposes any kinematic constraints applied to the musculoskeletal model. The cost function J of the musculoskeletal optimal control problem is a weighted sum of distinct pre-defined goals $G_i(\mathbf{x}, \mathbf{u})$. The weights and goals comprising the cost function are chosen based on the motion task. Example goals for a specific motion task (a sit-to-stand transfer) are given in Section III.

B. Upper-level Optimisation

Let $S = \{\mathbf{x}_i(t) \mid i = 1, 2, \dots, n\}$ be a dataset of state trajectories corresponding to a given motion. The aim of inverse optimal control is to identify the objective function J that best explains this set of motion samples. Assuming that the composition of goals in J is fixed, the mechanism for altering J is restricted to the suitable choice of the weights vector \mathbf{w} . Therefore, the optimisation of J can be expressed generically as follows in System 2:

$$\begin{aligned}
& \min_{\mathbf{w}} d(S, \mathbf{x}(t)), \\
& (\mathbf{x}(t), \mathbf{u}(t)) = F^P(\mathbf{w}), \\
& \mathbf{w} \in \mathbb{R}^n.
\end{aligned} \tag{2}$$

Here, we have encapsulated System 1 in the function F^P which maps weights \mathbf{w} on to optimal control solutions $(\mathbf{x}(t), \mathbf{u}(t))$, where P denotes the additional parameters that configure the forward optimal control problem (i.e. the bound and configuration vectors and underlying musculoskeletal model). The function d is a metric which quantifies the similarity or distance between the set of samples S and the state vector \mathbf{x} obtained at each iteration.

In practice, the optimisation problem expressed in System 2 is difficult to solve without additional constraints on the space of weights. Firstly, it is desirable to normalise the terms in the objective function such that similar weights lead to similar cost values. To achieve this, we define a set of basis weight vectors using standard basis vector notation, i.e. e_i denotes a weight vector with a 1 in the i^{th} coordinate and 0 elsewhere. Then, we compute the basis cost for each term of the objective function by running the forward optimal control problem on each normalisation weight vector. Mathematically, we compute a normalised cost matrix $C = \text{diag}(c_1, c_2, \dots, c_n)$ where:

$$c_i = \frac{1}{J(F^P(e_i))}. \tag{3}$$

Given C and an input weight vector \mathbf{w} , we compute the corresponding normalised weight vector $\bar{\mathbf{w}}$ as follows:

$$\bar{\mathbf{w}} = C\mathbf{w}. \tag{4}$$

Under this construction, a choice of comparable weights in the coordinates of \mathbf{w} translates to comparable constituent terms of the objective function $J(F^P(\bar{\mathbf{w}}))$.

¹OpenSim Moco is also capable of dealing with path constraints and time-invariant parameters, which are not relevant for our problem and so are excluded from our formulation.

In addition to normalisation, it is also useful to further constrain the space of weights to reduce redundancy. Consider that without such a constraint $\mathbf{w}_1 = (1, 1, \dots, 1)$ and $\mathbf{w}_2 = (2, 2, \dots, 2)$ are valid weight vectors, but from the perspective of relative cost ratios are functionally equivalent. To avoid such cases, we constrain the domain of the input weights, along with their sum, as follows:

$$\sum_{i=0}^n w_i = 1, \quad (5)$$

$$\mathbf{w} \in [0, 1]^n. \quad (6)$$

As a useful side-effect of this constraint and the normalisation procedure, a valid input weight vector can be interpreted as the relative importance of the terms in the objective function, e.g. $\mathbf{w} = (0.5, 0.3, 0.2)$ describes a situation in which G_1 , G_2 and G_3 contribute 50%, 30% and 20% to the total cost of the motion, respectively.

Incorporating the normalisation and non-redundancy constraints in to System 2 results in the following formulation for the upper-level optimisation:

$$\begin{aligned} & \min_{\mathbf{w}} d(S, \mathbf{x}(t)), \\ & (\mathbf{x}(t), \mathbf{u}(t)) = \mathbf{F}^P(\bar{\mathbf{w}}), \\ & \bar{\mathbf{w}} = C\mathbf{w}, \\ & C = \text{diag}(c_1, c_2, \dots, c_n), \\ & c_i = \frac{1}{J(\mathbf{F}^P(\mathbf{e}_i))}, \\ & \sum_{i=1}^n w_i = 1, \\ & \mathbf{w} \in [0, 1]^n. \end{aligned} \quad (7)$$

Similarly to the forward optimal control problem, we can encapsulate System 7 in the function I^P . The inverse musculoskeletal optimal control problem for a set of configuration parameters P and motion samples S can then be stated as:

$$\mathbf{w}^* = \arg \min I^P(S), \quad (8)$$

where \mathbf{w}^* is the optimal weights vector for tracking the state trajectories in S .

The optimiser used to solve Equation 8 is required to be gradient-free, support bounds and inequality constraints on parameters, and to perform well with computationally costly objective evaluations. In this work, we use NOMAD [20], [21], a black-box optimiser based on surrogate optimisation which performed favourably compared to Bayesian optimisation and genetic algorithm approaches in our initial testing. NOMAD does not support equality constraints, and so in practice Equation 5 was replaced with:

$$\sum_{i=1}^m w_i \leq 1, \quad (9)$$

for $m = n - 1$, and the following was prepended to the calculation of the objective function:

$$w_n = 1 - \sum_{i=1}^m w_i, \quad (10)$$

which has the favourable side effect of reducing the dimensionality of the problem by 1.

III. APPLICATION TO SIT-TO-STAND TRANSFER

When considering a specific motion task, we define an instantiation of Equation 8 by fixing the underlying musculoskeletal model, the objective function composition, and the task-specific constraints. Here, we consider inverse musculoskeletal optimal control applied to sit-to-stand transitions. We denote the sit-to-stand specific forward and inverse optimal control problems by F_{STS}^P and I_{STS}^P , respectively.

A. Musculoskeletal Model

The sit-to-stand biomechanics are modelled using a 2-dimensional musculoskeletal model with 4 joints representing the joint angles of the back, hip, knee and ankle. The joints are driven by coordinate actuators which respect activation dynamics. Contact forces, based on a smoothed Hunt-Crossley model, are used to model the physical interaction between the model and a seat, fixed at a specified height. Additional contact geometries are placed on the heel and toe bodies for the purpose of objective calculations. In Figure 1 the model is shown, alongside snapshots of a sit-to-stand movement.

B. Objective Composition

The set of goals comprising the objective function are a vital component of inverse optimal control problems. Ideally, the goal terms should be able to generate a rich variety of motions whilst being based on principles from movement biomechanics. In this work, we choose an objective function which consists of six goal terms. First, we include a measure of control effort:

$$g_1 = \frac{1}{d} \sum_{i=1}^4 u_i^3, \quad (\text{Effort})$$

where u_i is the i^{th} model control and d is the displacement of the centre of mass. Previous works have indicated that cubed controls are related to fatigue [22], [23].

Next, we include a measure of system stability, defined as:

$$g_2 = |p_{\text{com}} - p_{\text{bos}}|. \quad (\text{Stability})$$

Here, p_{com} is the extrapolated centre of mass, defined as:

$$p_{\text{com}} = x_{\text{com}} + v_{\text{com}} \sqrt{\frac{y_{\text{com}}}{g}}, \quad (11)$$

where x_{com} , y_{com} and v_{com} are the x position², y position and x velocity of the centre of mass, respectively, and g is acceleration due to gravity. The quantity p_{bos} is centre of the boundary of support in the x direction, measured as the midpoint between the heel and toe contact geometries. This stability goal is otherwise known as the margin of stability

²Here, and throughout this paper, we use the standard OpenSim coordinate system, where x is directed forwards from the point of view of the model, y is directed upwards, and z is directed to the right.

in the sagittal plane, which is related to dynamic stability when leaning [24].

Finally, the remaining goals compute the joint loading [25] at each degree of freedom of the model:

$$\begin{aligned} g_3 &= l_{\text{lumbar}}(t), && \text{(Lumbar Loading)} \\ g_4 &= l_{\text{hip}}(t), && \text{(Hip Loading)} \\ g_5 &= l_{\text{knee}}(t), && \text{(Knee Loading)} \\ g_6 &= l_{\text{ankle}}(t). && \text{(Ankle Loading)} \end{aligned}$$

The joint loading terms are calculated using OpenSim's underlying dynamics simulator Simbody [26].

Overall, using notation from System 1, the sit-to-stand specific objective function is as follows:

$$J_{\text{STS}} = \sum_{i=1}^6 w_i G_i. \quad (12)$$

This form of objective was chosen to balance the contributions from effort, stability, and joint-loading preferences that serve to personalise sit-to-stand behaviour for individual subjects.

C. Bounds and Constraints

The initial configuration of the model corresponds to a static sitting position, with $\mathbf{q}(0) = \mathbf{q}_s$ and $\dot{\mathbf{q}}(0) = \mathbf{0}$. The joint angle vector \mathbf{q}_s is typically obtained from individual subject measurements. The final configuration corresponds to a static standing position, i.e. $\mathbf{q}(T) = \dot{\mathbf{q}}(T) = \mathbf{0}$. Wide bounds on joint angles were chosen based on reference data in the literature [27], [28] as well as initial investigations on collected sit-to-stand data:

$$q_{\text{lumbar}} \in [-50, 15], \quad (13)$$

$$q_{\text{hip}} \in [-18, 135], \quad (14)$$

$$q_{\text{knee}} \in [-140, 5], \quad (15)$$

$$q_{\text{ankle}} \in [-60, 15], \quad (16)$$

while joint speeds were chosen to be expansive to allow for a wide range of viable motions:

$$\dot{\mathbf{q}} \in [-360, 360]^n. \quad (17)$$

Note that joint angles are specified in degrees and joint speeds are specified in degrees per second.

D. Upper-level Objective

The distance measure used to compare the forward optimal control motion with the set of sample trajectories is the mean continuous RMS of the joint angles, defined as follows:

$$d_{\text{STS}}(S, \mathbf{x}) = \frac{1}{n} \sum_{i=1}^n \sqrt{\frac{1}{mT} \int_0^T \left(\sum_{j=1}^m s_j^i(t) - x_j(t) \right)^2 dt}, \quad (18)$$

where S contains n motion samples which for convenience of notation we here write as $S = \{\mathbf{s}^1, \mathbf{s}^2, \dots, \mathbf{s}^n\}$. Without loss of generality we assume that the joint angles occupy positions 1 to m of the state vectors \mathbf{x} and \mathbf{s}^i .

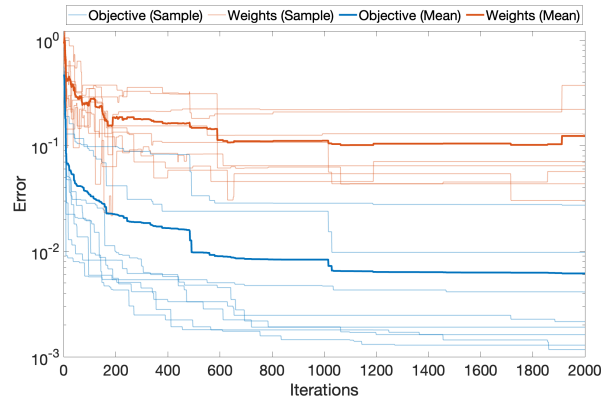


Fig. 2. The evolution of the objective RMSE (blue) and Euclidian distance error of the estimated weight vector (red) when solving for 8 pre-generated weights vectors, shown on a log scale. Each transparent line corresponds to a unique \mathbf{w}_{in} .

IV. COMPUTATIONAL VALIDATION

To evaluate the efficacy of the upper-level optimiser, we can first generate a solution to the forward problem using a known set of valid input weights, and directly compare those known weights to the result of the upper level optimiser. In other words, we execute the following sequence of operations:

- 1) Generate a valid input weight vector \mathbf{w}_{in} .
- 2) Compute $\bar{\mathbf{w}}_{\text{in}}$ and solve $F_{\text{STS}}^P(\bar{\mathbf{w}}_{\text{in}})$ for \mathbf{x}_{in} .
- 3) Set $S = \{\mathbf{x}_{\text{in}}\}$ and compute $\mathbf{w}_{\text{out}} = I_{\text{STS}}^P(S)$.
- 4) Evaluate $E_w = \|\mathbf{w}_{\text{in}} - \mathbf{w}_{\text{out}}\|_2$, $\mathbf{e} = \text{abs}(\mathbf{w}_{\text{in}} - \mathbf{w}_{\text{out}})$ and $E_F = |F_{\text{STS}}^P(\bar{\mathbf{w}}_{\text{in}}) - F_{\text{STS}}^P(\bar{\mathbf{w}}_{\text{out}})|$.

This process was carried out for a total of 8 inputs weights vectors, generated randomly subject to Equations 5 and 6. Graphs depicting the evolution of error in the objective function and weight vector estimate are shown in Figure 2. Table I reports the resulting errors E_F , E_w and \mathbf{e} averaged over all input weight vectors. Overall, the upper-level optimiser is able to obtain close approximations of the input weights, with a mean value of 0.123 for the error in the weight vector approximation. The weight estimates recovered had an error of 0.03 in the best case to 0.372 in the worst case.

V. EXPERIMENTAL VALIDATION

A. Evaluation of Perturbed Sit-to-Stand Biomechanics

The motion strategies employed by humans are highly context-dependant; for example, typical gait patterns are markedly different when walking on stable or unstable ground [29]. The same is true of sit-to-stand transfer, particularly when the motion becomes difficult due to underlying pathology or old age [30].

To investigate this phenomenon further, an experimental protocol was devised whereby 4 subjects were asked to complete a number of sit-to-stand transfers in the potential presence of perturbations, which induce instability in to the movement. The subjects carried out repetitive sit-to-stand transfers using a bench secured to a Motek Medical

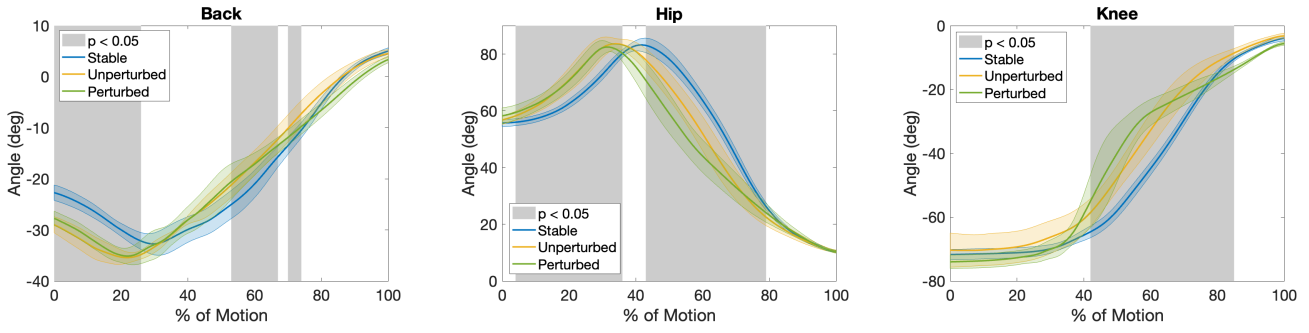


Fig. 3. A comparison of back, hip and knee joint angles during sit-to-stand transitions in the stable, unperturbed and perturbed modes for S1. Shaded regions indicate phases in which there was a statistically significant difference ($p < 0.05$) between the stable and unperturbed modes.

TABLE I
MEAN \pm STANDARD DEVIATION OF COMPUTATIONAL VALIDATION RESULTS

E_F	E_w	e_1	e_2	e_3	e_4	e_5	e_6
0.006 ± 0.009	0.123 ± 0.118	0.016 ± 0.029	0.014 ± 0.02	0.028 ± 0.038	0.042 ± 0.042	0.056 ± 0.067	0.082 ± 0.075

TABLE II
SIT-TO-STAND BIOMECHANICS EVALUATION

Subject	Mass (kg)	Range of Statistically Significant Difference (SSD) (% of Motion)				Total SSD (% of Motion)			
		Back	Hip	Knee	Ankle	Back	Hip	Knee	Ankle
1	92.5	0-26, 53-67, 70-74	4-36, 43-79	42-85	50-94	44	68	43	44
2	46.7	None	40-62	0-41, 50-60	0-60, 75-100	0	22	51	85
3	83.5	0-10, 99-100	0-16, 30-41	0-36, 98-100	0-51, 71-100	11	27	38	80
4	70.3	0-11, 31-100	0-8, 24-36, 49-100	52-79	27-54, 84-100	80	71	27	43

treadmill (Figure 1). A Vicon motion capture system was used to measure the kinematics of the back, hip, knee and ankle joints. Perturbations took the form of a brief negative acceleration of the treadmill belts, in order to simulate slippage of the feet in unstable conditions. The presence of a perturbation was determined randomly prior to each sit-to-stand transfer, with a likelihood of 50%. The perturbation was also randomly determined to begin at either 25%, 50% or 75% of max subject body weight being detected on the treadmill force plates. For reference, unperturbed sit-to-stand data was also collected for each subject, with the likelihood of perturbation set to 0%.

Following data collection, the resulting marker data was filtered, segmented and then processed via inverse kinematics in OpenSim to compute joint angle trajectories. The resulting motions were classified in to one of three categories:

- *Stable*, in which the perturbation likelihood was 0%.
- *Unperturbed*, in which there was a 50% likelihood of perturbation, but no perturbation occurred.
- *Perturbed*, in which there was a 50% likelihood of perturbation, and a perturbation did occur.

In Figure 3, we show sit-to-stand kinematics of the back, hip, and knee joints indexed by sit-to-stand perturbation mode for one subject as a visual example. In particular, the regions of the sit-to-stand transfer at which there is a statistically

significant difference between stable and unperturbed biomechanics are indicated via shading. Statistical significance was calculated using Welch's t -test. The full data on statistically significant ranges for combination of subject and perturbation mode is provided in Table II. Each subject contributed 5 motion samples to each mode.

B. Sit-to-Stand Strategy Modulation

The results presented in Table II show that, even when a perturbation does not occur, subject awareness of the potential perturbation induces a statistically significant difference in sit-to-stand biomechanics. This suggests that humans modulate their sit-to-stand motion strategy in the presence of potential instabilities.

To test this hypothesis, we used the inverse musculoskeletal optimal control framework to compute the weights vector for each subject across both the stable and unperturbed sit-to-stand modes, via the following sequence of steps:

- 1) Set S_s and S_u to be the sets of motion samples corresponding to the stable and unperturbed modes.
- 2) Compute $\mathbf{w}_s = I_{STS}^P(S_s)$ and $\mathbf{w}_u = I_{STS}^P(S_u)$.
- 3) Evaluate $|\mathbf{w}_s - \mathbf{w}_u|$ and $|d_{STS}(S_s, \mathbf{x}_s) - d_{STS}(S_u, \mathbf{x}_u)|$, which quantify the difference in underlying motion strategy and in tracking performance, respectively.

In Figure 4 we compare for visualisation purposes the resultant stable and unperturbed weight vectors, along with

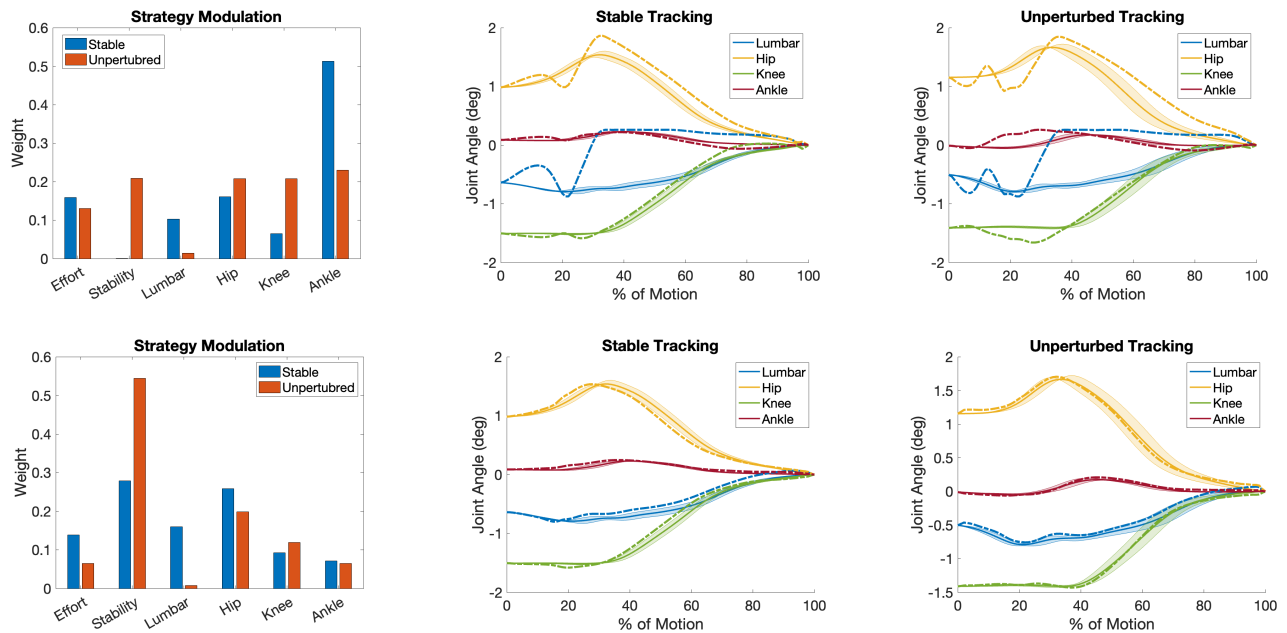


Fig. 4. **Top row:** analysis of strategy modulation between the stable and unperturbed sit-to-stand modes for Subject 3. Additional plots are shown which indicate the tracking performance of the forward optimal control problem using the recovered weights. Solid lines show the mean of experimentally measured kinematic trajectories, shaded regions show their standard deviation, and dash-dot lines show the predicted motion. **Bottom row:** a similar analysis, but where the ankle loading goal has been replaced by a tracking goal. This improves the tracking performance of the lumbar joint.

TABLE III
SIT-TO-STAND STRATEGY MODULATION RESULTS

Subject	Unperturbed Weights						Perturbed Weights						Tracking Error (d_{STS})	
	w_1	w_2	w_3	w_4	w_5	w_6	w_1	w_2	w_3	w_4	w_5	w_6	Stable	Unperturbed
1	0.17	0	0.151	0.217	0.168	0.294	0.293	0	0.2	0	0.168	0.339	0.178	0.221
2	0.016	0.045	0.363	0.064	0.083	0.429	0	0.367	0.067	0	0.267	0.3	0.134	0.260
3	0.159	0	0.103	0.161	0.065	0.513	0.131	0.209	0.014	0.208	0.208	0.230	0.235	0.194

the associated tracking errors, for Subject 3. As a point of comparison, we also show the results obtained when replacing the ankle loading goal g_6 with a tracking goal. The full results for strategy modulation and tracking performance for each subject are provided in Table III.

VI. DISCUSSION

The computational validation in Section IV illustrates that the inverse musculoskeletal optimal control method presented in this work is able to recover close approximations to known input weights. In addition, the experimental results from Section V show firstly that humans modulate their underlying sit-to-stand strategy in the presence of perturbations (Figure 3) and furthermore that the nature of this modulation is highly subject-specific (Table II). For example, Subject 2 largely modulates knee and ankle angle trajectories in the presence of potential disturbances, while Subject 4 largely modulates back and hip angle trajectories.

Similarly, we see that the recovered motion strategies for each subject (Table III) are also subject-specific. Subjects 2 and 3 both increase the weight of the stability goal at the cost of loading terms, while Subject 1 increases the lumbar

loading term, with no influence from stability. Notably, the tracking error obtained when running the framework on experimental data (Table III) is higher than when using simulated data (Figure 2). In particular, the deviations in lumbar angle for Subject 3 are not recreated when tracking is included in the cost function (Figure 4). This suggests that the components of the overall cost function $J(\mathbf{x}, \mathbf{u})$ do not fully span the set of experimental sit-to-stand trajectories.

A key benefit of this framework lies in the generation of human strategy models which can potentially be used as human motion predictors for human-robot collaboration. As evidenced in Figure 4, the performance of these models is likely strongly linked to the choice of the underlying objective composition used in the IOC process. In the future, we intend to explore a wider range of cost function compositions which can be directly compared in their ability to track experimental data, thus providing biomechanical insights into the criteria underpinning human sit-to-stand strategies. We also intend to explore the generalisability of our framework to other motion tasks. Moving forward, we will quantify the predictive power of the obtained models in experiments with human-exoskeleton systems.

REFERENCES

- [1] V. Villani, F. Pini, F. Leali, and C. Secchi, "Survey on human-robot collaboration in industrial settings: Safety, intuitive interfaces and applications," *Mechatronics*, vol. 55, pp. 248–266, 2018.
- [2] J. P. Pinho, C. Taira, P. Parik-Americano, L. O. Suplino, V. P. Bartholomeu, V. N. Hartmann, G. S. Umemura, and A. Forner-Cordero, "A comparison between three commercially available exoskeletons in the automotive industry: An electromyographic pilot study," in *2020 8th IEEE RAS/EMBS International Conference for Biomedical Robotics and Biomechatronics (BioRob)*. IEEE, 2020, pp. 246–251.
- [3] A. Plaza, M. Hernandez, G. Puyuelo, E. Garces, and E. Garcia, "Lower-limb medical and rehabilitation exoskeletons: A review of the current designs," *IEEE Reviews in Biomedical Engineering*, 2021.
- [4] G. S. Sawicki, O. N. Beck, I. Kang, and A. J. Young, "The exoskeleton expansion: improving walking and running economy," *Journal of neuroengineering and rehabilitation*, vol. 17, no. 1, pp. 1–9, 2020.
- [5] T. Stouraitis, I. Chatzinikolaidis, M. Gienger, and S. Vijayakumar, "Online hybrid motion planning for dyadic collaborative manipulation via bilevel optimization," *IEEE Transactions on Robotics*, vol. 36, no. 5, pp. 1452–1471, 2020.
- [6] 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.
- [7] D. F. Gordon, C. McGreavy, A. Christou, and S. Vijayakumar, "Human-in-the-loop optimization of exoskeleton assistance via online simulation of metabolic cost," *IEEE Transactions on Robotics*, vol. 38, no. 3, pp. 1410–1429, 2022.
- [8] S. O'Connor, "Exoskeletons in nursing and healthcare: a bionic future," *Clinical Nursing Research*, vol. 30, no. 8, pp. 1123–1126, 2021.
- [9] M. Geravand, P. Z. Korondi, and A. Peer, "Human sit-to-stand transfer modeling for optimal control of assistive robots," in *5th IEEE RAS/EMBS International Conference on Biomedical Robotics and Biomechatronics*. IEEE, 2014, pp. 670–676.
- [10] R. Zaman, Y. Xiang, R. Rakshit, and J. Yang, "Hybrid predictive model for lifting by integrating skeletal motion prediction with an opensim musculoskeletal model," *IEEE Transactions on Biomedical Engineering*, vol. 69, no. 3, pp. 1111–1122, 2021.
- [11] B. Q. Hoa, V. Padois, F. Benamar, and E. Desailly, "A two-step optimization-based synthesis of squat movements," in *International Conference on Human-Computer Interaction*. Springer, 2021, pp. 122–138.
- [12] V. Q. Nguyen, R. T. Johnson, F. C. Sup, and B. R. Umberger, "Bilevel optimization for cost function determination in dynamic simulation of human gait," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 27, no. 7, pp. 1426–1435, 2019.
- [13] K. Mombaur, A. Truong, and J.-P. Laumond, "From human to humanoid locomotion—an inverse optimal control approach," *Autonomous robots*, vol. 28, no. 3, pp. 369–383, 2010.
- [14] H. El-Hussieny, A. Asker, and O. Salah, "Learning the sit-to-stand human behavior: an inverse optimal control approach," in *2017 13th international computer engineering conference (ICENCO)*. IEEE, 2017, pp. 112–117.
- [15] M. Geravand, P. Z. Korondi, C. Werner, K. Hauer, and A. Peer, "Human sit-to-stand transfer modeling towards intuitive and biologically-inspired robot assistance," *Autonomous Robots*, vol. 41, no. 3, pp. 575–592, 2017.
- [16] R. M. Alexander, *Optima for animals*. Princeton University Press, 1996.
- [17] A. Alamdari and V. N. Krovi, "A review of computational musculoskeletal analysis of human lower extremities," *Human modelling for bio-inspired robotics*, pp. 37–73, 2017.
- [18] A. Seth, J. L. Hicks, T. K. Uchida, A. Habib, C. L. Dembia, J. J. Dunne, C. F. Ong, M. S. DeMers, A. Rajagopal, M. Millard *et al.*, "Opensim: Simulating musculoskeletal dynamics and neuromuscular control to study human and animal movement," *PLoS computational biology*, vol. 14, no. 7, p. e1006223, 2018.
- [19] C. L. Dembia, N. A. Bianco, A. Falisse, J. L. Hicks, and S. L. Delp, "Opensim moco: Musculoskeletal optimal control," *PLOS Computational Biology*, vol. 16, no. 12, p. e1008493, 2020.
- [20] S. Le Digabel, "Algorithm 909: Nomad: Nonlinear optimization with the mads algorithm," *ACM Transactions on Mathematical Software (TOMS)*, vol. 37, no. 4, pp. 1–15, 2011.
- [21] C. Audet, S. L. Digabel, V. R. Montplaisir, and C. Tribes, "Nomad version 4: Nonlinear optimization with the mads algorithm," *arXiv preprint arXiv:2104.11627*, 2021.
- [22] M. Ackermann and A. J. Van den Bogert, "Optimality principles for model-based prediction of human gait," *Journal of biomechanics*, vol. 43, no. 6, pp. 1055–1060, 2010.
- [23] R. D. Crowninshield and R. A. Brand, "A physiologically based criterion of muscle force prediction in locomotion," *Journal of biomechanics*, vol. 14, no. 11, pp. 793–801, 1981.
- [24] A. Hof, M. Gazendam, and W. Sinke, "The condition for dynamic stability," *Journal of biomechanics*, vol. 38, no. 1, pp. 1–8, 2005.
- [25] K. M. Steele, M. S. DeMers, M. H. Schwartz, and S. L. Delp, "Compressive tibiofemoral force during crouch gait," *Gait & posture*, vol. 35, no. 4, pp. 556–560, 2012.
- [26] M. A. Sherman, A. Seth, and S. L. Delp, "Simbody: multibody dynamics for biomedical research," *Procedia Iutam*, vol. 2, pp. 241–261, 2011.
- [27] J. Soucie, C. Wang, A. Forsyth, S. Funk, M. Denny, K. Roach, D. Boone, and H. T. C. Network, "Range of motion measurements: reference values and a database for comparison studies," *Haemophilia*, vol. 17, no. 3, pp. 500–507, 2011.
- [28] J. E. Bible, D. Biswas, C. P. Miller, P. G. Whang, and J. N. Grauer, "Normal functional range of motion of the lumbar spine during 15 activities of daily living," *Clinical Spine Surgery*, vol. 23, no. 2, pp. 106–112, 2010.
- [29] G. Martino, Y. P. Ivanenko, A. d'Avella, M. Serrao, A. Ranavolo, F. Draicchio, G. Cappellini, C. Casali, and F. Lacquaniti, "Neuromuscular adjustments of gait associated with unstable conditions," *Journal of neurophysiology*, vol. 114, no. 5, pp. 2867–2882, 2015.
- [30] P. O. Riley, D. E. Krebs, and R. A. Popat, "Biomechanical analysis of failed sit-to-stand," *IEEE Transactions on rehabilitation engineering*, vol. 5, no. 4, pp. 353–359, 1997.