Adaptive Reference Inverse Optimal Control for Natural Walking With Musculoskeletal Models

Jiacheng Weng, Ehsan Hashemi, *Member, IEEE*, and Arash Arami, *Member, IEEE*

***Abstract*— An efficient inverse optimal control method named Adaptive Reference IOC is introduced to study natural walking with musculoskeletal models. Adaptive Reference IOC utilizes efficient inner-loop direct collocation for optimal trajectory prediction along with a gradient-based weight update inspired by structured classification in the outer-loop to achieve about 7 times faster convergence than existing derivative-free methods while maintaining similar outcomes in terms of gait trajectory matching. The proposed method adequately reconstructed the reference data when applied to experimental walking data from ten participants walking at various speeds and stride lengths. The proposed framework can facilitate efficient personalized cost function optimization for specific walking tasks, and provide guidance to personalized reference trajectory design for assistive robotic systems such as lower-limb exoskeletons.**

***Index Terms*— Direct collocation, gait, inverse optimal control, musculoskeletal model, predictive simulation, structured prediction.**

I. Introduction

HUMAN locomotion is the result of a complex interaction between the central nervous system (CNS), peripheral nervous system, musculoskeletal system, and the environment. The recent development of open-source software tools with detailed and accurate musculoskeletal models allows researchers to simulate human natural walking with ease [1], [2]. Assuming the natural walking behaviour results from some optimization done by our CNS in interaction with the environment [3], researchers have generated realistic dynamic walking by formulating it as an optimal control problem with a set of neuromechanically sound cost terms [1], [4]. The optimal control problem is then solved using trajectory optimization techniques such as direct collocation (DC) or learning-based methods such as reinforcement learning (RL).

Manuscript received November 26, 2021; revised March 18, 2022 and May 5, 2022; accepted May 31, 2022. Date of publication June 6, 2022; date of current version June 14, 2022. This work was supported in part by NSERC Discovery under Grant RGPIN-2018-04850 and Grant RGPIN-2020-05097, in part by the of John R. Evans Leaders Fund Canadian Foundation for Innovation, in part by the Ontario Research Fund (ORF), and in part by the New Frontiers in Research Fund under Grant NFRFE-2018-01698. *(Corresponding author: Arash Arami.)*

Jiacheng Weng is with the Mechanical and Mechatronics Engineering Department, University of Waterloo, Waterloo, ON N2L 3G1, Canada (e-mail: j5weng@uwaterloo.ca).

Ehsan Hashemi is with the Mechanical Engineering Department, University of Alberta, Edmonton, T6G 1H9, Canada (e-mail: ehashemi@ualberta.ca).

Arash Arami is with the Mechanical and Mechatronics Engineering Department, University of Waterloo, Waterloo, ON N2L 3G1, Canada, and also with the Toronto Rehabilitation Institute (KITE), University Health Network, Toronto, M5G 2C4, Canada (e-mail: arash.arami@uwaterloo.ca).

This article has supplementary downloadable material available at https://doi.org/10.1109/TNSRE.2022.3180690, provided by the authors.

Digital Object Identifier 10.1109/TNSRE.2022.3180690

One common challenge of the optimal control problems for human locomotion is to find the optimal cost function. Researchers often used a linear combination of multiple cost terms that captures energy expenditure, muscle activation, motion smoothness, and stability [1], [4]–[9], but the weights associated with the cost terms were often manually tuned, which is slow and labour-intensive. Recently, Weng *et al.* developed a strategic weight update using curriculum learning [9] for cost function adaptation. Their cost weight update method was implemented in a deep RL agent that efficiently learns the muscle excitations needed for walking [9]. Some studies utilized global optimization techniques for weight tuning but were computationally intensive. For example, the genetic algorithm (GA) from Nguyen *et al.* [7] took 139 hours on a 10-core i9 CPU to optimize the cost function weights. This limitation is the main reason why these studies solve only the cost function tuning problem for one specific setup without analysis of the cost function generalization over different body shapes and locomotion task parameters.

Inverse optimal control (IOC) is a class of methods that aim to recover optimal cost functions given observed state trajectories. Common methods include inverse Karush–Kuhn–Tucker (KKT) [10], [11], structured classification (SC) [12], [13], hyperparameter tuning [14], and inverse reinforcement learning (IRL) [15], [16]. Some of these methods utilize gradient-based updates providing higher efficiency than derivative-free optimization techniques. However, these methods have not yet been closely examined in challenging locomotion studies considering musculoskeletal models.

Herein, we propose our algorithm called Adaptive Reference IOC to discover the optimal cost function of walking in detailed musculoskeletal models. The Adaptive Reference IOC utilizes DC as the inner-loop optimal controller, and gradient-based weight updates in the outer loop to achieve high efficiency during the weight update process. The main contributions of the paper are summarized as:

- formulate an efficient IOC framework for natural walking in detailed musculoskeletal models with a significant computational efficiency improvement compared to the GA
- demonstrate the success of the proposed method in recovering cost functions based on multiple participant-specific reference data

II. Related Work

In this section, we discuss various cost function tuning methods, including KKT conditions, derivative-free optimization methods, SC, hyperparameter tuning techniques,

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

and IRL. Their use cases and limitations for locomotion are also discussed.

A. Inverse Optimal Control

IOC addresses the problem of recovering the underlying cost function given the desired state trajectories such as human demonstrations. IOC problems are often formulated as bilevel optimization problems with a parametric cost function given a priori. The inner-loop problems find optimal state trajectories given the cost function, whereas the outer-loop problems tune the parameters of the cost function so that the inner-loop solutions resemble the desired state trajectories. For problems that involve only kinematics such as motion path prediction (e.g., obstacle avoidance during motion), KKT conditions have been imposed to optimize the weights with high computational efficiency (e.g., the inverse KKT method in [10] optimized cost function weights 144 times faster than derivative-free optimization for robotic arm manipulation; [11] used KKT conditions to optimize cost functions for locomotion path with sub-second run time). However, for more complex cost functions that involve system dynamics (e.g., joint torques), derivative-free optimization techniques such as COBYLA (for humanoid [17]) and CMA-ES (for robotic arms [18]) are used. These techniques are based on strategic cost weight sampling coupled with evaluations (i.e., solving optimization problems) of all sampled cost functions. However, for locomotion problems that utilize musculoskeletal models, the internal optimization problems (usually in the form of trajectory optimization problems) are more challenging to solve than torque-driven humanoid and robotic arms due to redundant muscle actuators and complex muscle dynamics. In the existing locomotion studies with musculoskeletal models, the cost function weights are often manually tuned based on the researcher's domain knowledge [1], [4], [8]. With a reasonable simplification of the musculoskeletal models by constraining sagittal-plane motion and reducing the number of muscles, researchers formulated the IOC problem for natural walking using GA for optimal cost weight search [7], however it still requires a long computational time (139 hours on a 10-core i9 CPU). On the other hand, SC treats the IOC problem as a structured prediction problem (i.e., a multi-class classification problem with an extremely large number of classes) and allows for a gradient-based update to the cost function weights, which can be more efficient than derivative-free methods [12], [13]. Direct loss minimization (DLM), which is the successor to the vanilla structured prediction, achieves better performance by using a loss-adjusted cost function for lower-level optimization (i.e., adding the domain-specific loss term to push the optimized trajectories towards the preferred region of search) during weight update [19]. DLM's performance marks it as a good candidate for solving the IOC for complex locomotion problems.

B. Inverse Reinforcement Learning

IRL is a specific branch of IOC that solves expert policies (i.e., a closed-loop controller capable of sampling all optimal trajectories given initial conditions) in the inner-loop optimization instead of specific trajectories. Maximum margin IRL solves the optimal cost function by increasing the distance between expert policies and all other policies [13], [20], [21]. Maximum entropy IRL [15], [16], on the other hand, finds the optimal cost function that matches the expectation of the expert behaviour while maximizing the entropy of the policy. All of these IRL approaches require solving a forward RL problem (i.e., learning an expert policy through the agent's interaction with the environment) either fully or partially which is non-trivial by itself. This makes the extension of these methods to continuous and high-dimensional settings (e.g., muscle actuated human locomotion) difficult [22], [23].

In some applications (e.g., computer graphics) where the control policy is the primary need, imitation learning methods are used to solve the IRL problem where the cost function tuning is absorbed into the policy optimization [24], [25]. These methods provide little insight into important factors that shape human locomotion and therefore are unsuitable for this study.

C. Hyperparameter Optimization

IOC problems with parametric cost functions can also be seen as a hyperparameter tuning problem [26] commonly found in the machine learning field [14]. While naive optimization methods, including grid search and random search [27], can be used for IOC with easy parallelization, they are rarely applied to problems with many parameters (weights). This is due to an exponential increase in the number of weight samples required to be evaluated in order to find close-to-optimal weights. Sequential model-based optimization (SMBO) is an alternative approach that utilizes a surrogate model (i.e., an approximation model that mimics the observed input-output relationship; often obtained through data-driven approaches) to estimate the relationship between the total cost values and the cost function weights [28]. Bayesian optimization is a common SMBO method that uses probabilistic models (e.g., Gaussian processes) to locate the good weight candidates during the weight search [29]. Similar to other derivative-free methods mentioned in II-A, SMBO may be slow for human locomotion study as the inner-loop optimization problem is difficult to solve.

III. Methods

We first discuss the customization of the subject-specific musculoskeletal models, followed by the formulation of the inner-loop optimization using DC, and last the outer-loop Adaptive Reference IOC algorithm inspired by SC and DLM [12], [19]. In this study, a publicly available dataset for treadmill walking [30] is used for the IOC analysis. The dataset includes ten healthy participants walking at various speeds and stride lengths.

A. Musculoskeletal Models

Natural gait patterns vary across individuals due to their diverse body geometries, muscle fibre type proportions, strength, and possible variations in neural control. When performing cost function optimization for individual participants, a personalized musculoskeletal model is required to better

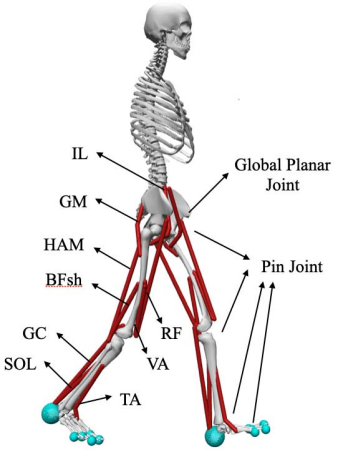

Fig. 1. The musculoskeletal model has 11 DoFs and 18 muscles.

represent the physical characteristics. Here, we first construct a baseline musculoskeletal model, then scale it to match the anthropometric parameters of each participant.

The baseline musculoskeletal model is customized based on the *gait10dof18musc* model from OpenSim [31] which is shown in Fig. 1. The model allows sagittal-plane (2D) motion with 11 degrees of freedom (DoF), including 8 DoF at the hip, knee, ankle, and metatarsophalangeal (toe) joints, and 3 DoF for global translation and rotation at the pelvis. The full system configuration is described using a 40-dimensional state vector, including the position and velocity of all 11 DoFs and 18 muscle activation signals. The 18-dimensional control signals specify muscle excitation, which is linked to muscle activation using first-order activation dynamics [31]. Lower limb muscles include hamstrings (HAM), biceps femoris short head (BFsh), gluteus maximus (GM), iliopsoas (IL), rectus femoris (RF), vastus (VA), gastrocnemius (GC), soleus (SOL), and tibialis anterior (TA). The hip, knee, and ankle joints are softly constrained using passive joint torque actuators with double exponent formulation [32] to simulate the effect of ligaments. The metatarsophalangeal joints are passively actuated using linear passive spring-damper actuators. The default muscle fibre lengths and pennation angles are obtained from [33]. The default tendon slack lengths are obtained from [34]. Six Hunt Crossley contact spheres at the heel, first metatarsal head (MH), third MH, fifth MH, hallux, and middle toe are added for each foot to cover all high-pressure zones during walking [35].

The subject-specific musculoskeletal models are obtained by scaling the baseline model. The body geometries, contact sphere locations, and optimal muscle fibre lengths are scaled based on the marker tracking data. The model mass is also scaled to the reported subject's body mass while maintaining the same mass distribution as the baseline model. The maximum isometric muscle force is calculated using the method explained in [34]. First, the subject's body mass and height are taken as input to estimate the individual muscle volumes using the experimental results from [36]. The muscle volumes are then used along with the scaled optimal muscle fibre lengths and the muscle-specific tension from [37] to estimate the maximum muscle isometric forces.

B. Inner-Loop Optimization

The inner-loop trajectory optimization assumes a given cost function and finds the optimal trajectory that minimizes the given cost using DC. We used OpenSim Moco [2] to formulate the DC problem with musculoskeletal models. The inner-loop cost function for generating the walking motion consists of six terms with four variable weights:

$$\begin{aligned} J_{inner} &= \frac{1}{d}\int_0^{t_f} (\boldsymbol{\omega}^T \boldsymbol{C}) dt \\ &= \frac{1}{d}\int_0^{t_f} \Bigg[10 \sum_i \left(\frac{u_i^3}{n}\right) + 10^{-2} \sum_j \hat{\tau}_j^2 \\ &\quad + \omega_1 \times 10(\theta_{torso})^2 \\ &\quad + \omega_2 \times \left(\hat{a}_{CoM}\right)^2 \\ &\quad + \omega_3 \times 10^2 \left[\left(\hat{F}^y_{lf} \hat{v}^{fx}_{lf}\right)^2 + \left(\hat{F}^y_{rf} \hat{v}^{fx}_{rf}\right)^2 \right] \\ &\quad + \omega_4 \times 10(\hat{p}_{CoMx} - \hat{p}_{midBoS})^2 \Bigg] dt, \end{aligned} \tag{1}$$

which can be seen as the integral of a dot product of the weight vector $\boldsymbol{\omega}$ and the cost vector $\boldsymbol{C}$. As experimental data from multiple participants is used in our study, non-dimensionless quantities including torque, acceleration, force, position, and velocity exhibit different ranges across the participants due to factors such as segment length and body mass (can be seen in Table II). These quantities are scaled to the equivalent dimensionless forms [38] using the following equations:

$$\hat{\tau} = \frac{\tau}{mgl_{leg}}, \quad \hat{F} = \frac{F}{mg}, \quad \hat{v} = \frac{v}{\sqrt{l_{leg} g}}, \quad \hat{p} = \frac{p}{l_{leg}}, \tag{2}$$

where m is the subject mass, l_{leg} is the leg length, g is $9.81 m/s^2$. Such scaling improves the cost function generalization across different body sizes. The first term penalizes the cubed muscle excitation u where i is the muscle index, and n is the total number of muscles. This term has been widely used in previous studies to capture the CNS effort to minimize energy expenditure and muscle fatigue [1], [4], [5], [7], [8]. The second term penalizes the use of scaled passive ligament torque $\hat{\tau}_j$ similar to [1], [4], [8] to keep joint angles within the limits of the physiological range of motion, where j is the index of the passive joint actuators at hip, knee, and ankle joints. The third term penalizes the deviation of torso angle θ_{torso} from the upright position similar to [7] as an approximation of trunk energy expenditure. The fourth term penalizes the scaled acceleration of the center of mass (CoM) inspired by [1] to encourage motion smoothness. The fifth term penalizes the foot sliding similar to [6] to ensure proper foot clearance. This is evaluated based on the scaled vertical ground reaction force (GRF) $\hat{F}^y$ and the scaled horizontal foot velocity $\hat{v}^{fx}$, where lf, rf represent left and right feet. The last term encourages the stability (similar to [7]) by penalizing deviation of the scaled forward CoM position $\hat{p}_{CoMx}$ from the center of the extended base of support (BoS) $\hat{p}_{midBoS}$ (i.e., the mid-point between the toes of the leading foot and the heel of the back foot projected onto the ground plane). Coefficients of the cost terms scale the cost values to the same order of magnitude given the desired motion

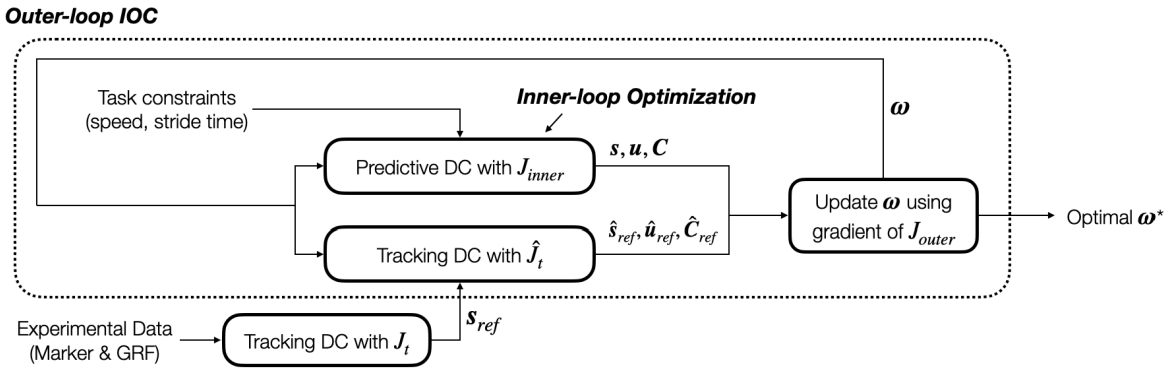

Fig. 2. Block diagram of the Adaptive Reference IOC algorithm. The tracking DC problem outside of the outer-loop IOC module generates feasible trajectories s_{ref}. The inner-loop predictive DC problem generates state trajectories s, control trajectories u, and the cost vector C subject to stride time and average speed constraints, as well as state and control bounds. The tracking DC problem inside the outer-loop IOC module provides adaptive reference trajectories for the weight update. All DC problems are solved with the scaled musculoskeletal models using methods described in III-A.

trajectories from the experiment. After integration, the final cost is divided by the total forward distance d measured at the CoM. The first two terms in the cost function are designed with fixed constant weights due to two reasons. First, uniformly scaling the multipliers of all cost terms does not change the optimal solution. To prevent uniform scaling, the muscle activation multiplier is set to a constant so that the weight updates ensure relative weight changes between cost terms and, consequently, changes in the solution. Second, the passive joint torque term is fixed to improve the stability of the optimization. During natural walking, the lower-limb joints rarely reach extreme angles resulting in small cost values (i.e., close to zero). This can cause the weight update to the passive joint torque term to be unreliable due to poor numerical precision. In some cases, the weight may converge to zero, which may cause the inner-loop optimization to generate abnormal walking with knee hyperextension and subsequently lead to a large rebound of the cost weight. Therefore, the weight of passive joint torque term is fixed to prevent its convergence to zero and ensure the stability of the optimization program.

To reduce the complexity of the DC problem, the optimized walking trajectory is limited to half a gait cycle with gait symmetry constraints. These symmetry constraints ensure mirrored state values (i.e., position and velocity of all DoFs except the global forward distance, and all muscle activation signals) about the sagittal plane at the trajectory start and end. After solving the DC problem, the full gait cycle is reconstructed by concatenating the trajectories from the left and the right leg. Several equality and inequality constraints are imposed on the state and control variables. Muscle excitation and activation are constrained to $[0, 1]$. Joint angles and global body positions are constrained to feasible ranges (i.e., $[-15°, 60°]$, $[-70°, 5°]$, $[-15°, 25°]$, for hip, knee, and ankle angles respectively; [0m, 1m], [0.75m, 1.25m], $[-20°, 20°]$ for global forward distance, height, and orientation at pelvis respectively). The average walking speed and the final trajectory time are constrained to the experimental values to maintain the same walking task. Last, the trajectories are discretized to only 15 nodes which are shown to be sufficient by [7] and our investigations.

Algorithm 1 Adaptive Reference IOC

Initialize: desired trajectories s_k^{ref}, s_f^{ref} from solving DC problems using J_t; initial weight vector ω; learning rate η

while *max iter. not reached* **do**

$\hat{s}_{ref}, \hat{u}_{ref} = \underset{s,u}{\text{argmin}} \hat{J}_t(\omega)$;

compute cost vector $\hat{C}_{ref} = C(\hat{s}_{ref}, \hat{u}_{ref})$;

$s', u' = \underset{s,u}{\text{argmin}} J_{inner}(\omega)$;

compute cost vector $C' = C(s', u')$;

$\omega \leftarrow \omega + \eta(C' - \hat{C}_{ref})$;

end

C. Outer-Loop Optimization

The outer-loop optimization aims to optimize ω in the cost function J_{inner} such that the inner-loop solution matches the reference trajectories from the experiments. With the inspiration from SC [12] and DLM [19], we formulate the Adaptive Reference IOC for the outer-loop optimization shown in Alg. 1 and Fig. 2. Adaptive Reference IOC consists of three main steps. First, the experimental gait data is collected and pre-processed to ensure dynamic feasibility. Next, two DC problems are solved in parallel with one generating predictive gait trajectories given the current cost weights, and the other one generating adaptive reference trajectories by adding experimental data tracking to the current cost function. Last, the converged trajectories from these two DC problems are compared to update the cost weights. Steps 2 and 3 are repeated until the convergence criteria are satisfied.

We first represent the walking task-related parameters (e.g., walking speed and stride time) as x, and the walking trajectories (including both states and controls) as y. If the cost weight ω is optimal (i.e., $\omega = \omega^*$), then the cost of y_{ref} as the dot product of ω and cost vector C is smaller than all other choices of $y \in \mathcal{Y}$, where $\mathcal{Y}$ is the superset of all possible trajectories. such a relationship can be represented as:

$$\omega^T C(x, y_{ref}(x)) \leq \omega^T C(x, y(x)) + \zeta, \tag{3}$$

where $\zeta \geq 0$ deals with possible violation of the inequality when $\boldsymbol{\omega}$ is sub-optimal. Intuitively, we want ζ to be small to ensure the optimality of $\boldsymbol{\omega}$. In addition, the following inequality holds for any arbitrary $\boldsymbol{\omega}$:

$$\boldsymbol{\omega}^T \boldsymbol{C}(\boldsymbol{x}, \boldsymbol{y}_{ref}(\boldsymbol{x})) \geq \min_{\boldsymbol{y}\in\mathcal{Y}} \boldsymbol{\omega}^T \boldsymbol{C}(\boldsymbol{x}, \boldsymbol{y}(\boldsymbol{x})). \tag{4}$$

The inequality in (4) reduces to equality when $\boldsymbol{\omega}$ is optimal (i.e., $\boldsymbol{\omega} = \boldsymbol{\omega}^*$). As (3) holds for all $\boldsymbol{y}$, we can substitute $\boldsymbol{\omega}^T \boldsymbol{C}(\boldsymbol{x}, \boldsymbol{y}(\boldsymbol{x}))$ in (3) by its minimum over $\boldsymbol{y}$. We then obtain:

$$\zeta \geq \boldsymbol{\omega}^T \boldsymbol{C}(\boldsymbol{x}, \boldsymbol{y}_{ref}(\boldsymbol{x})) - \min_{\boldsymbol{y}\in\mathcal{Y}} \boldsymbol{\omega}^T \boldsymbol{C}(\boldsymbol{x}, \boldsymbol{y}(\boldsymbol{x})) \geq 0. \tag{5}$$

Since we want small ζ, we have to minimize its lower bound. Therefore, the problem becomes:

$$\min_{\boldsymbol{\omega}} J_{outer} = \min_{\boldsymbol{\omega}} [(\boldsymbol{\omega}^T \boldsymbol{C}(\boldsymbol{x}, \boldsymbol{y}_{ref}(\boldsymbol{x})) - \min_{\boldsymbol{y}\in\mathcal{Y}} \boldsymbol{\omega}^T \boldsymbol{C}(\boldsymbol{x}, \boldsymbol{y}(\boldsymbol{x}))]. \tag{6}$$

The gradient of the objective function with respect to $\boldsymbol{\omega}$ can then be calculated as follows:

$$\nabla J_{outer} = \boldsymbol{C}(\boldsymbol{x}, \boldsymbol{y}_{ref}) - \boldsymbol{C}(\boldsymbol{x}, \boldsymbol{y}') \tag{7}$$

$$\boldsymbol{y}' = \operatorname*{argmin}_{\boldsymbol{y}\in\mathcal{Y}} \boldsymbol{\omega}^T \boldsymbol{C}(\boldsymbol{x}, \boldsymbol{y}(\boldsymbol{x})), \tag{8}$$

where $\boldsymbol{y}_{ref}$ is the reference trajectory from experiments, $\boldsymbol{y}'$ is the optimal inner-loop DC solution given $\boldsymbol{\omega}$. As $\boldsymbol{C}(\boldsymbol{x}, \boldsymbol{y}_{ref})$ relies on clean control signals (muscle excitation) which are nearly impossible to collect during experiments, we estimate the control signals by solving a tracking DC problem using the cost function $\hat{J}_t$:

$$\hat{J}_t = J_{inner} + \frac{1}{d}\int_0^{t_f} c_t dt \tag{9}$$

$$c_t = \omega_k \|\boldsymbol{s}_k - \boldsymbol{s}_k^{ref}\|_2^2 + \omega_f \|(\boldsymbol{s}_f - \boldsymbol{s}_f^{ref})\|_2^2, \tag{10}$$

where $\boldsymbol{s}_k$, $\boldsymbol{s}_f$ are the kinematic and GRF trajectories to be optimized in this tracking DC problem, and $\boldsymbol{s}_k^{ref}$, $\boldsymbol{s}_f^{ref}$ are the reference trajectories from experiment. Weights ω_k and ω_f control the trade-off between the kinematic and the kinetic tracking. As J_{inner} is a function of $\boldsymbol{\omega}$, $\hat{J}_t$ varies during the weight optimization. This requires the tracking DC problem to be re-solved after each weight update. Intuitively, we allow the estimated reference trajectory $\boldsymbol{y}_{ref}$ to adapt to the optimized weight vector while moving towards the reference motion. This modification nicely results in the DLM for structured prediction which was originally proposed in [19]. Consequently, we replace $\boldsymbol{C}(\boldsymbol{x}, \boldsymbol{y}_{ref})$ in (7) with $\boldsymbol{C}(\boldsymbol{x}, \hat{\boldsymbol{y}}_{ref})$ where $\hat{\boldsymbol{y}}_{ref}$ is the adaptive reference generated by the tracking DC problem using (9). Then gradient descent is used for the weight update. Comparing to using a static reference $\boldsymbol{y}_{ref}$ (further away from $\boldsymbol{y}'$), the adaptive reference $\hat{\boldsymbol{y}}_{ref}$ (closer to $\boldsymbol{y}'$) also reduces the inaccuracy in calculating the cost value differences in (7) due to numerical errors from the direct collocation problem (e.g., trajectory reconstruction errors), and improve the stability of the cost update. The resulting algorithm for learning optimal weights (that indicates the optimal cost function) for natural walking, called Adaptive Reference IOC, is summarized in Alg. 1, and Fig. 2 shows the connection among different components of this algorithm.

In actual experiments, the experimental data $\boldsymbol{s}_k^{ref}$, $\boldsymbol{s}_f^{ref}$ in (10) is pre-processed by solving a tracking DC problem with the cost function J_t:

$$J_t = \frac{1}{d}\int_0^{t_f} \left(c_t + \|\boldsymbol{u}\|_2^2\right) dt. \tag{11}$$

This is to convert treadmill walking to level ground walking in simulation and ensure dynamic feasibility of the reference trajectory. The simple gradient descent update rule is replaced with the momentum update to accelerate convergence. The learning rate η is tuned using grid search during the experiment with synthetic data (explained in section IV-A) to achieve stable and efficient convergence. The tuned η is then applied to all other experiments. A dynamically adjusted weight lower bound (set to 1/10 of the initial weights, and reduced by half if the bound is activated for more than 5 consecutive iterations) and gradient clipping are also used to avoid the weight speedy drop to zero due to the momentum update which can cause significantly different inner-loop DC solutions and outer-loop weight oscillations.

IV. Results and Discussion

In this section, we first test the convergence of the Adaptive Reference IOC algorithm using synthetic data generated from DC with known weight compositions. Then, the algorithm is applied to treadmill walking data of ten subjects. The performance and limitations of the IOC method are also discussed. Last, we compare the performance of the Adaptive Reference IOC with GA over all ten participants. The ω_k and ω_f in (10) are set to 0.01 and 0.005 for all experiments to allow deviations from the tracking data. Such deviation allows the adaptive reference trajectory to be closer to the trajectories obtained using the current cost weights, which avoids large gradient updates and smooths out the gradient history. The difference between ω_k and ω_f ensures an approximately similar magnitude of the kinematic tracking cost and the ground reaction tracking cost for subjects.

A. IOC Performance on Synthetic Data

To test the convergence of the proposed algorithm, we used a synthetic reference trajectory by solving a DC problem with known weights in (1). The DC problem used the scaled musculoskeletal model for the first participant in the dataset [30]. The walking task is configured to be at 1.2m/s with a stride frequency of 1Hz. The synthetic known weights in (1) are all configured to 0.5 for demonstration. To test the robustness of the Adaptive Reference IOC algorithm against the weight initialization, we used three different sets of initial weights to represent different initialization conditions (i.e., smaller / larger than the known weights, and mix of both) as shown in table I. All three trials were run in parallel and converged after 100 weight updates which took about 3 hours of total computation using AMD Ryzen 7 3700X CPU at 3.6GHz. The final converged weights and the weight errors are shown in Table I and Fig. 3, respectively. The final converged weights are similar to the desired values for all three trials, and the weight error progress showed a nice reduction over the

TABLE I
WEIGHT COMPARISON AFTER SOLVING THE SYNTHETIC IOC PROBLEM

	Initial Weights $[\omega_1, \omega_2, \omega_3, \omega_4]$	IOC solutions $[\omega_1, \omega_2, \omega_3, \omega_4]$
Trial 1	[2.5, 2.5, 2.5, 2.5]	[0.57, 0.49, 0.54, 0.46]
Trial 2	[0.1, 0.1, 0.1, 0.1]	[0.56, 0.48, 0.55, 0.46]
Trial 3	[0.1, 0.1, 2.5, 2.5]	[0.57, 0.49, 0.55, 0.46]
Desired weights		[0.50, 0.50, 0.50, 0.50]

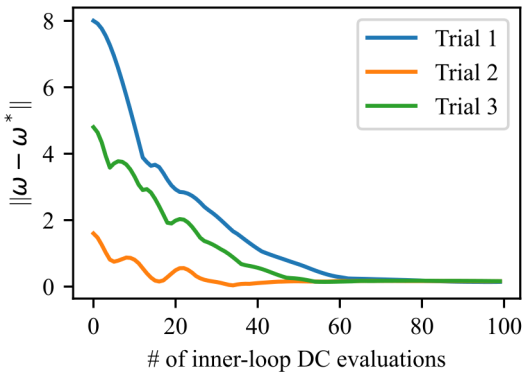

Fig. 3. Weight vector error reduction plot using synthetic reference trajectories. Initial and final weights for all three trials are listed in Table I.

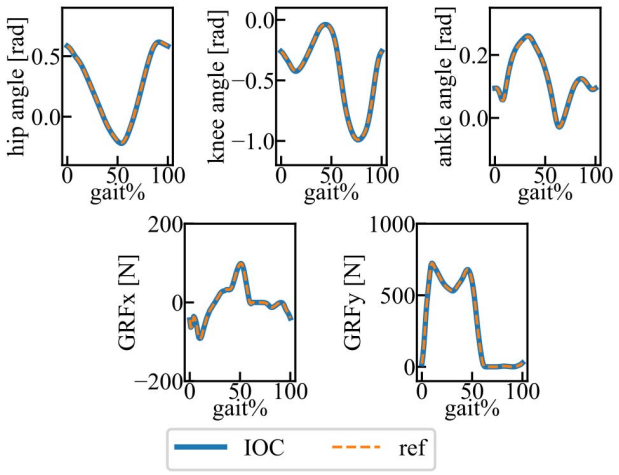

Fig. 4. Comparison of gait kinematics and GRFs between the IOC solution and the synthetic reference gait.

inner-loop evaluations. Trial 1 and 2 required more inner-loop updates as the initial weights are further away from the reference weights. Some oscillations exist in the weight error plot which are primarily caused by the weight overshoot from the momentum update. The final weight errors stabilized slightly above 0 after iteration 80 for all three trials. This is because the adaptive reference trajectories $\hat{\boldsymbol{s}}_{ref}, \hat{\boldsymbol{u}}_{ref}$ in Fig. 2 allows small deviation from the reference trajectories in exchange for a slightly easier outer-loop optimization problem. This deviation is small and can be neglected as shown in Fig. 4.

B. IOC Performance on Experimental Data

In this section, treadmill walking data of ten healthy participants [30] were used. The initial weight vector $[\omega_1, \omega_2, \omega_3, \omega_4]$ is set to [1, 1, 1, 1] for all subjects. The

TABLE II
SUBJECT MASS, HEIGHT, LEG LENGTH, SPEED, STRIDE TIME AND OPTIMAL WEIGHTS BY SOLVING IOC PROBLEMS

#	M kg	h m	l_{leg} m	v m/s	t_s s	ω_1	ω_2	ω_3	ω_4
1	64.9	1.71	0.87	1.31	1.08	0.112	0.241	1.634	0.589
2	57.7	1.58	0.82	1.26	0.99	0.010	0.329	1.638	0.606
3	70.7	1.77	0.94	1.05	1.30	0.521	0.142	1.466	0.088
4	70.2	1.77	0.95	1.10	1.18	0.001	0.514	0.757	0.202
5	79.2	1.79	0.97	1.05	1.17	0.027	0.787	1.180	0.330
6	73.5	1.78	0.98	1.16	1.22	0.714	0.584	0.717	0.445
7	73.0	1.61	0.86	0.95	1.06	0.788	0.845	1.116	0.226
8	73.1	1.76	0.93	1.05	1.26	0.988	1.077	0.958	0.324
9	94.3	1.85	1.02	1.00	1.25	0.001	1.045	1.833	0.160
10	59.5	1.54	0.79	1.10	1.08	0.010	0.488	0.452	0.404

IOC problems were solved with 193 weight updates on average (i.e., about 10 hours of computation per problem using AMD Ryzen 7 3700X CPU at 3.6GHz). Comparing to other derivative-free methods for cost function determination with similar setup which solved 5600 inner-loop optimization problems [7], the Adaptive Reference IOC significantly reduces the number of inner-loop evaluations (385 problems on average shown in Table IV). The subject information, walking configurations, and the optimized weights are summarized in Table II. The converged CoM acceleration weights ω_2 (0.605 ± 0.325, mean$\pm$standard deviation (SD) computed across subjects), foot sliding weights ω_3 (1.175 ± 0.460), and stability weights ω_4 (0.337 ± 0.175) are considerably larger than zero for all participants. This behaviour can indicate the significance of these cost terms in capturing natural walking behaviour. In contrast, the torso orientation weights ω_1 for five participants (participant 2, 4, 5, 9, and 10) converge to small values (either 0.001, 0.01, or 0.027). This may be due to poor estimation of the desired torso orientation in the reference trajectories. As the motion tracking data only includes lower limbs, the estimation of the torso orientation solely depends on reconstructed pelvis orientation, which can have excessive forward tilt due to marker placement errors. Such inaccuracy leads to a large torso orientation cost for the reference trajectories, which essentially pushes ω_1 to zero. In addition, potential redundancy in the cost function may also cause some cost weights to converge to zero.

Given the optimized weights in Table II, the optimal joint kinematic and GRF trajectories are computed and compared to the reference trajectories. Fig. 5 shows the gait comparison of a typical participant (participant 1 in [30]). Overall, the joint angles and GRFs from IOC match closely with the tracking solution. The ankle joint angles from IOC have a smoother profile than the reference during the stance phase.

Fig. 6 shows the muscle activation plots obtained from the optimal cost function (through IOC) of the same participant. The EMG signals of four lower-limb muscles, which were measured and available in [30], are overlaid in Fig. 6. The EMG signals are scaled to best match the magnitude of the muscle activation from the optimal cost function for visualization. The EMG signals of HAM, GC, and TA showed good alignment to the obtained muscle activation trajectories.

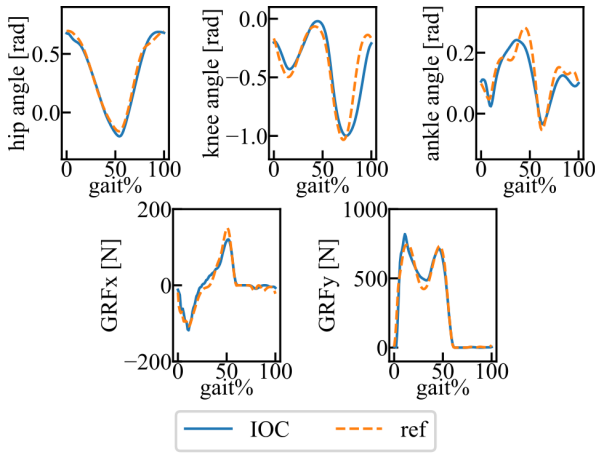

Fig. 5. Comparison of joint angles and GRFs between the IOC solution and the experimental reference gait from the first participant in [30]; GRFx and GRFy are along the forward and vertical directions respectively.

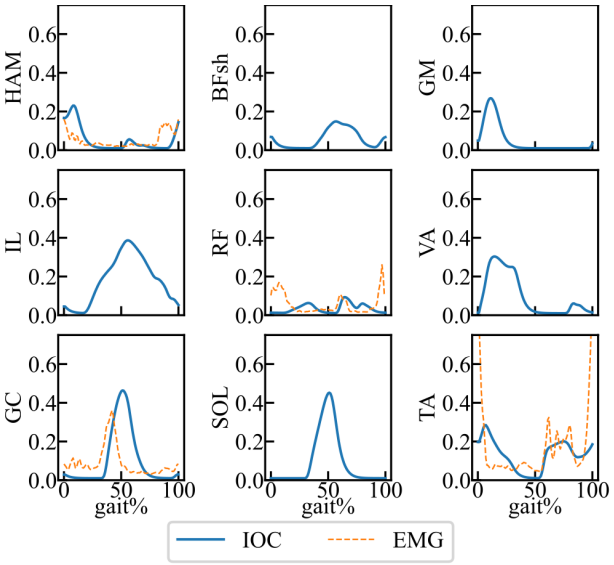

Fig. 6. Comparison of muscle activations from IOC and average EMG signals of four measured muscles of the first participant in [30]. EMG signals for the other five muscles were not measured in [30].

The EMG signal of RF shows good alignment to the combined muscle activation from two knee extensors (i.e., RF and VA). There may be two explanations for the aforementioned observations. First, surface EMG sensors only measure lumped muscle action potential at the location of measurement which can cause inaccuracies in individual muscle measurement when sensor placement is sub-optimal. Second, the individual muscle weights in the muscle activation penalty term may not be optimal, which may favour activating VA over RF. The obtained GC activation is slightly delayed compared to the EMG signal similar to [7]. This may be caused by the problem reduction from 3D to the sagittal plane where delayed toe push-off is required to facilitate the delayed knee extension for proper foot clearance.

To analyze the gait pattern quality, gait errors are computed with respect to the subject-specific reference trajectories using root mean square errors (RMSE). After IOC weight optimization, hip, knee and ankle RMSEs of $0.041 \pm 0.007, 0.122 \pm 0.017, 0.054 \pm 0.009$ in radians, and GRFx

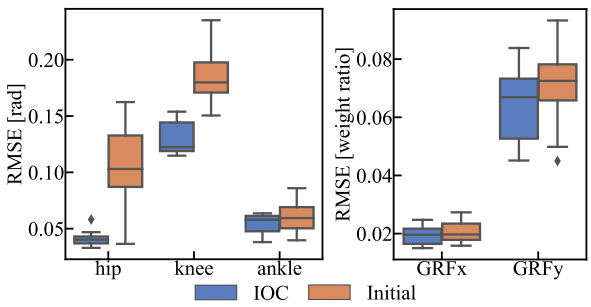

Fig. 7. RMSE of the joint angles and GRFs before and after solving the IOC problems; GRFx and GRFy are along the forward and vertical directions respectively.

and GRFy RMSEs of $0.020 \pm 0.004, 0.064 \pm 0.013$ in terms of body weight ratio (i.e., GRF error [N] / body weight [N]) were obtained, respectively. The hip and ankle joint angles generally align well with the reference after weight optimization. The knee angles have relatively larger errors due to the delayed knee extension (visible in Fig. 5). The ground reaction forces also align well with the reference. RMSEs of the trajectories obtained from the initial weights are also computed to confirm the gait improvement from the weight optimization, which is summarized in Fig. 7. Such comparison indicates that the improvement mainly happens in kinematics features, especially hip and knee angles. GRF profiles start with small errors and show mild improvement from the weight modification. The small initial errors for the GRF are potentially due to the constraints on the task parameters (i.e., walking speed and stride length) during the IOC experiments (walking speed and stride length can impact GRF features like peak height and stance duration). In addition, the optimized gait patterns have Pearson correlations of 0.992 ± 0.004, 0.917 ± 0.024, 0.843 ± 0.077 for hip, knee, ankle angles, and 0.957 ± 0.024, 0.990 ± 0.003 for GRFx and GRFy respectively.

Last, a correlation analysis is performed between the optimized weights and parameters, including body mass, walking speed, and stride time (raw data available in Table II). As the CNS weighting of different cost terms may vary for different participants (e.g. mass, height, muscle properties) and walking tasks (e.g. speed, stride time), this correlation analysis could bring insights into how the cost weights vary as a function of those parameters. This is valuable particularly when we need to infer an optimal cost function for a walking task when we do not have access to its specific reference data. As the correlation is calculated across participants that have various body geometries, the walking speed and the stride time are scaled to their dimensionless forms using the following equations from [38]:

$$\hat{v} = \frac{v}{\sqrt{l_{leg}\, g}}, \quad \hat{t}_s = \frac{t_s}{\sqrt{l_{leg}/g}} \tag{12}$$

where l_{leg} is the leg length, g is $9.81 m/s^2$. The Pearson correlation coefficients are summarized in Table III. In terms of participant parameters, the body mass is negatively correlated to the scaled walking speed meaning slower walking for participants with larger body mass. The stride time is also negatively correlated to the scaled walking speed. In terms of weights, the torso orientation weights ω_1 and the foot sliding weights ω_3 show low correlations to the resultant

TABLE III
PEARSON CORRELATION COEFFICIENTS BETWEEN SCALED WALKING PARAMETERS AND WEIGHTS

	M	$\hat{v}$	$\hat{t}_s$	ω_1	ω_2	ω_3	ω_4
M	1	-0.78	0.34	0.07	0.68	0.31	-0.60
$\hat{v}$	-0.78	1	-0.52	-0.35	-0.67	0.10	0.88
$\hat{t}_s$	0.34	-0.52	1	0.43	0.13	-0.17	-0.63

scaled walking speed and stride time. The CoM acceleration weight ω_2 has a negative correlation with the scaled walking speed suggesting less demand on the smooth motion at higher speeds. The stability weight has a positive correlation with the scaled walking speed, indicating more penalty to stability with faster walking. However, such an observation is not sufficient to draw a concrete conclusion as the correlation is computed based on only ten participants walking with limited gait speed and stride length variations. Extending this study to include variable gait data of each participant, considered as future work, would allow probing the existence of such speed-cost relationships.

One main challenge in using the proposed IOC framework to solve optimal locomotion cost weights is the stability of the inner-loop DC problems. DC transcripts the trajectory optimization problem (i.e., in the form of differential equations) into constrained parameter optimization problems with algebraic equations by discretizing the trajectories and reconstructing them using function approximations. The discretized optimization program is then solved by a non-linear solver such as IPOPT [39]. Such an optimization technique does not guarantee convergence to the global optima and is often sensitive to the initial trajectory guesses, especially for high-dimensional complex systems such as musculoskeletal models. We found that DC problems sometimes converge poorly if the initial trajectory guesses are infeasible (e.g. the musculoskeletal model sinks into the solid ground) or far off from the walking motion (e.g. jumping motion is given with two legs moving together which also satisfies symmetry constraints). We also observed that the optimized trajectories alternate between two configurations during IOC (e.g. with and without extensive knee extension during the stance phase). This indicates that the slightly changing weights cause the DC problem to converge to different local optima, which can cause oscillations in the outer-loop weight updates. This issue was mitigated by imposing proper coordinate constraints as explained in III-B, fixing the weight to the passive joint torque term in J_{inner}, and using initial trajectories with feasible walking motion. However, it is impossible to consider all edge cases that can lead to inner-loop DC instability. In addition, improving the expressive capacity of the defined cost function may improve the inner-loop DC stability. This improvement essentially worsens the other local optima and expand the convex region around the desired motion.

C. IOC-GA Performance Comparison

To compare the Adaptive Reference IOC algorithm with derivative-free optimization methods in optimizing cost weights for natural walking, we solve the IOC problems for all ten participants using the GA implementation in MATLAB R2021a and compare the performance in Table IV. The bounds for weight sampling are set to $[0, 4]$. A population with a size of 40 is generated with random initial weights uniformly sampled within the weight bounds. The fitness function for the GA explicitly minimizes the sum of tracking errors for hip angle, knee angle, ankle angle, GRFx and GRFy. The converged weights from GA are generally larger than the ones from Adaptive Reference IOC except for ω_1 (Adaptive Reference IOC exhibits significantly larger ω_1 for subjects 6, 7, and 8). When excluding these three subjects, the Pearson correlation between the obtained weight vectors is 0.930 ± 0.058 (0.731 ± 0.329 when including all subjects). The high Pearson correlation coefficients for seven subjects indicate that the ratio between the four cost weights contributes more to the final converged gait than the weight magnitudes. The small differences in joint angle errors and Pearson correlations between the Adaptive Reference IOC and GA for subjects 6, 7, 8 suggest that our cost function formulation may have redundancy which allows similar converged gait with different cost weights.

TABLE IV
COMPARISON BETWEEN ADAPTIVE REFERENCE IOC AND GA

	Adaptive Reference IOC (mean±SD)	GA (mean±SD)
ω_1	0.317 ± 0.392	0.076 ± 0.068
ω_2	0.605 ± 0.325	1.200 ± 0.918
ω_3	1.175 ± 0.460	2.142 ± 0.893
ω_4	0.337 ± 0.175	0.707 ± 0.290
e_{hip}[rad]	0.041 ± 0.007	0.042 ± 0.020
e_{knee}[rad]	0.122 ± 0.017	0.114 ± 0.010
e_{ankle}[rad]	0.054 ± 0.009	0.056 ± 0.010
e_{GRFx} [weight ratio]	0.020 ± 0.004	0.018 ± 0.010
e_{GRFy} [weight ratio]	0.064 ± 0.013	0.064 ± 0.010
r_{hip}	0.992 ± 0.004	0.993 ± 0.002
r_{knee}	0.917 ± 0.024	0.938 ± 0.011
r_{ankle}	0.843 ± 0.077	0.815 ± 0.070
r_{GRFx}	0.957 ± 0.024	0.965 ± 0.021
r_{GRFy}	0.990 ± 0.003	0.991 ± 0.003
# of DC eval.	385 ± 157	3124 ± 941
computing time [hr]	10 ± 4	69 ± 21

Despite the differences in weights, both methods produce gait patterns similar to the reference trajectories with no statistical differences for most errors and correlation measures between the Adaptive Reference IOC and GA ($p > 0.05$ using the Wilcoxon Signed Rank test) except r_{knee} (GA outperforms Adaptive Reference IOC with $p = 0.006$) and r_{ankle} (Adaptive Reference IOC outperforms GA with $p = 0.012$). Although e_{GRFx} and r_{GRFx} showed statistical differences ($p = 0.001$ for both cases), such differences can be neglected as the mean difference is small (i.e., less than 1%). The subtle differences in converged gait between GA and the Adaptive Reference IOC can be caused by their different ways of gait evaluation as GA considers joint angle and GRF errors directly, whereas the Adaptive Reference IOC considers the cost values which is the abstraction of the original gait. The Adaptive Reference IOC only requires about 12% of the internal DC evaluations and 14% of the computing time compared to the GA. Besides the computational improvement,

the proposed algorithm, unlike GA, has no requirement on the initial bounds for weight vector sampling.

V. Conclusion

An efficient IOC algorithm named Adaptive Reference IOC is proposed for the natural walking problem with detailed musculoskeletal models. We showcased the efficiency of the proposed algorithm in tuning cost functions and matching gait trajectories using both synthetic data and experimental data. The optimized cost weights and the corresponding optimal gait trajectories from our algorithm were also compared with the solutions from GA. With the Adaptive Reference IOC, preliminary findings about the correlation between optimal weights and walking task parameters (i.e., walking speed and stride time) were reported. We also discussed the limitation of the proposed algorithm in terms of stability of the inner-loop DC problems and limitation of the experimental data. In future studies, we plan to enhance the existing cost function analysis by enriching the experimental data (i.e., including various walking speeds, various stride lengths, upper body kinematic tracking, and more muscles and EMG measurement sites). The proposed method can be used to assist exoskeleton controller design by providing personalized cost function and reference trajectory tuning.

References

[1] A. Falisse, G. Serrancolí, C. L. Dembia, J. Gillis, I. Jonkers, and F. De Groote, "Rapid predictive simulations with complex musculoskeletal models suggest that diverse healthy and pathological human gaits can emerge from similar control strategies," *J. Roy. Soc. Interface*, vol. 16, no. 157, Aug. 2019, Art. no. 20190402.

[2] C. L. Dembia, N. A. Bianco, A. Falisse, J. L. Hicks, and S. L. Delp, "OpenSim MOCO: Musculoskeletal optimal control," *PLOS Comput. Biol.*, vol. 16, no. 12, Dec. 2020, Art. no. e1008493.

[3] E. Todorov and M. I. Jordan, "Optimal feedback control as a theory of motor coordination," *Nature Neurosci.*, vol. 5, no. 11, pp. 1226–1235, 2002.

[4] B. R. Umberger, "Stance and swing phase costs in human walking," *J. Roy. Soc. Interface*, vol. 7, no. 50, pp. 1329–1340, Sep. 2010.

[5] M. Ackermann and A. J. van den Bogert, "Optimality principles for model-based prediction of human gait," *J. Biomech.*, vol. 43, no. 6, pp. 1055–1060, 2010.

[6] I. Mordatch, J. M. Wang, E. Todorov, and V. Koltun, "Animating human lower limbs using contact-invariant optimization," *ACM Trans. Graph.*, vol. 32, no. 6, pp. 1–8, Nov. 2013.

[7] 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 Trans. Neural Syst. Rehabil. Eng.*, vol. 27, no. 7, pp. 1426–1435, Jul. 2019.

[8] A. Falisse *et al.*, "Physics-based simulations to predict the differential effects of motor control and musculoskeletal deficits on gait dysfunction in cerebral palsy: A retrospective case study," *Frontiers Hum. Neurosci.*, vol. 14, p. 40, Feb. 2020.

[9] J. Weng, E. Hashemi, and A. Arami, "Natural walking with musculoskeletal models using deep reinforcement learning," *IEEE Robot. Autom. Lett.*, vol. 6, no. 2, pp. 4156–4162, Apr. 2021.

[10] P. Englert, N. A. Vien, and M. Toussaint, "Inverse KKT: Learning cost functions of manipulation tasks from demonstrations," *Int. J. Robot. Res.*, vol. 36, nos. 13–14, pp. 1474–1488, 2017.

[11] A.-S. Puydupin-Jamin, M. Johnson, and T. Bretl, "A convex approach to inverse optimal control and its application to modeling human locomotion," in *Proc. IEEE Int. Conf. Robot. Automat.*, May 2012, pp. 531–536.

[12] N. D. Ratliff, J. A. Bagnell, and M. A. Zinkevich, "(Approximate) Subgradient methods for structured prediction," in *Proc. 7th Int. Conf. Artif. Intell. Statist.*, 2007, pp. 380–387.

[13] N. D. Ratliff, J. A. Bagnell, and M. A. Zinkevich, "Maximum margin planning," in *Proc. 23rd Int. Conf. Mach. Learn. (ICML)*, 2006, pp. 729–736.

[14] M. Feurer and F. Hutter, "Hyperparameter optimization," in *Automated Machine Learning*. Cham, Switzerland: Springer, 2019, pp. 3–33.

[15] B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey, "Maximum entropy inverse reinforcement learning," in *Proc. AAAI*, Chicago, IL, USA, vol. 8, 2008, pp. 1433–1438.

[16] C. Finn, S. Levine, and P. Abbeel, "Guided cost learning: Deep inverse optimal control via policy optimization," in *Proc. 33rd Int. Conf. Mach. Learn.*, 2016, pp. 49–58.

[17] D. Clever, R. M. Schemschat, M. L. Felis, and K. Mombaur, "Inverse optimal control based identification of optimality criteria in whole-body human walking on level ground," in *Proc. 6th IEEE Int. Conf. Biomed. Robot. Biomech. (BioRob)*, Jun. 2016, pp. 1192–1199.

[18] A. Doerr, N. Ratliff, J. Bohg, M. Toussaint, and S. Schaal, "Direct loss minimization inverse optimal control," in *Proc. Robot., Sci. Syst.*, Italy, Rome, Jul. 2015.

[19] D. A. McAllester, T. Hazan, and J. Keshet, "Direct loss minimization for structured prediction," in *Proc. Adv. Neural Inf. Process. Syst.*, Princeton, NJ, USA, 2010, vol. 1, no. 2, p. 3.

[20] P. Abbeel and A. Y. Ng, "Apprenticeship learning via inverse reinforcement learning," in *Proc. 21st Int. Conf. Mach. Learn. (ICML)*, 2004, pp. 1–8.

[21] N. Ratliff, J. A. Bagnell, and S. S. Srinivasa, "Imitation learning for locomotion and manipulation," in *Proc. 7th IEEE-RAS Int. Conf. Humanoid Robots*, Nov. 2007, pp. 392–397.

[22] Ł. Kidziński *et al.*, "Learning to run challenge solutions: Adapting reinforcement learning methods for neuromusculoskeletal environments," in *The NIPS'17 Competition: Building Intelligent Systems*. Cham, Switzerland: Springer, 2018, pp. 121–153.

[23] Ł. Kidzinski *et al.*, "Artificial intelligence for prosthetics: Challenge solutions," in *The NeurIPS'18 Competition: From Machine Learning to Intelligent Conversations*. Cham, Switzerland: Springer, 2019, pp. 69–128.

[24] J. Ho and S. Ermon, "Generative adversarial imitation learning," 2016, *arXiv:1606.03476*.

[25] A. Boularias, J. Kober, and J. Peters, "Relative entropy inverse reinforcement learning," in *Proc. 14th Int. Conf. Artif. Intell. Statist.*, 2011, pp. 182–189.

[26] J. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl, "Algorithms for hyper-parameter optimization," in *Proc. 25th Adv. Neural Inf. Process. Syst.*, vol. 24, 2011, pp. 1–9.

[27] J. Bergstra and Y. Bengio, "Random search for hyper-parameter optimization," *J. Mach. Learn. Res.*, vol. 13, no. 2, pp. 1–25, 2012.

[28] F. Hutter, H. H. Hoos, and K. Leyton-Brown, "Sequential model-based optimization for general algorithm configuration," in *Proc. Int. Conf. Learn. Intell. Optim.* Cham, Switzerland: Springer, 2011, pp. 507–523.

[29] J. T. Wilson, F. Hutter, and M. P. Deisenroth, "Maximizing acquisition functions for Bayesian optimization," 2018, *arXiv:1805.10196*.

[30] R. Macaluso, K. Embry, D. Villarreal, and R. Gregg. (2020). *Human Leg Kinematics, Kinetics, and EMG During Phase-Shifting Perturbations at Varying Inclines*. [Online]. Available: https://dx.doi.org/10.21227/12hp-e249

[31] S. L. Delp *et al.*, "OpenSim: Open-source software to create and analyze dynamic simulations of movement," *IEEE Trans. Biomed. Eng.*, vol. 54, no. 11, pp. 1940–1950, Nov. 2007.

[32] F. C. Anderson and M. G. Pandy, "Dynamic optimization of human walking," *J. Biomech. Eng.*, vol. 123, no. 5, pp. 381–390, Oct. 2001.

[33] S. R. Ward, C. M. Eng, L. H. Smallwood, and R. L. Lieber, "Are current measurements of lower extremity muscle architecture accurate?" *Clin. Orthopaedics Rel. Res.*, vol. 467, no. 4, pp. 1074–1082, 2009.

[34] V. Nguyen, "Predictive simulation of human movement and applications to assistive device design and control," Ph.D. dissertation, Mech. Eng., Univ. Massachusetts Amherst, Amherst, MA, USA, 2019.

[35] T. C. Pataky and J. Y. Goulermas, "Pedobarographic statistical parametric mapping (pSPM): A pixel-level approach to foot pressure image analysis," *J. Biomech.*, vol. 41, no. 10, pp. 2136–2143, Jul. 2008.

[36] G. Handsfield, C. H. Meyer, J. Hart, M. Abel, and S. S. Blemker, "Relationships of 35 lower limb muscles to height and body mass quantified using MRI," *J. Biomech.*, vol. 47, no. 3, pp. 631–638, 2014.

[37] A. Rajagopal, C. L. Dembia, M. S. DeMers, D. D. Delp, J. L. Hicks, and S. L. Delp, "Full-body musculoskeletal model for muscle-driven simulation of human gait," *IEEE Trans. Biomed. Eng.*, vol. 63, no. 10, pp. 2068–2079, Oct. 2016.

[38] A. L. Hof, "Scaling gait data to body size," *Gait, Posture*, vol. 3, no. 4, pp. 222–223, 1996.

[39] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Math. Program.*, vol. 106, no. 1, pp. 25–57, May 2006.