Wearable Robotics Assisted Bipedal Balance and Locomotion Control with Whole-body Musculoskeletal Simulations

Anonymous Author(s)
Affiliation
Address
email

Abstract:
The design and control of wearable robotics to assist human movement present a significant challenge, largely due to the complexity of human-robot interaction. Traditional development cycles that rely on extensive hardware prototyping and human subject testing are often costly, time-consuming, and pose potential safety risks. We present a novel simulation framework for the rapid, real-time optimization of wearable robotics hardware and control parameters. Our approach leverages a detailed, whole-body musculoskeletal model to simulate dynamic interactions between the human user and an exoskeleton during both locomotion and balance tasks. This collaborative simulation enables the efficient evaluation of various wearable robotics configurations and control strategies, providing deep insights into their effects on human biomechanics, including muscle activation patterns and stability. We demonstrate the framework's efficacy by optimizing hip exoskeletons for assistance, proving that this approach can rapidly provide actionable insights into device design and control, thereby accelerating the development and reducing the risks associated with creating assistive robotic technologies.

Keywords: wearable robotics, balance control, human locomotion, musculoskeletal system

1 Introduction

Bipedal locomotion and balance have been extensively studied in robotics and control, with research demonstrating the inherent challenges of stabilizing underactuated, high-degree-of-freedom systems [1, 2, 3]. In humanoid robotics, stability is often analyzed using criteria such as limit cycles and gait periodicity [4, 5, 6]. However, these challenges are not exclusive to bipedal humanoid robots. Maintaining stable bipedal posture and gait remains a complex task for humans, for these functions rely on precise neuromuscular coordination, multi-sensory integration, and the nonlinear dynamics of the human musculoskeletal system.

Recent advances in musculoskeletal simulation have enabled biomechanically realistic models of human movement, providing unique opportunities to study balance, locomotion, and muscle-level mechanics in simulation [7, 8, 9, 10, 11]. Recent work introduced a full-body human musculoskeletal system that simulates whole-body dynamics [12]. Such models offer controlled, repeatable conditions that are difficult or unsafe to achieve in physical experiments, making them well suited for investigating interactions between humans and assistive devices under various conditions. This capability is particularly relevant for wearable robotics such as exoskeletons, whose influence on human balance, gait, and muscle activation is often difficult to quantify comprehensively in live trials due to hardware limitations, safety concerns, and measurement constraints.

Submitted to the Human to Robot (H2R) Workshop at 9th Conference on Robot Learning (CoRL 2025). Do not distribute.

In this work, we extend full-body musculoskeletal simulation to develop a validation platform for exoskeleton–human interaction. Our framework enables rapid evaluation of how wearable robotics morphology and assistance control policies affect human gait, stability, and muscle activation demands. By integrating detailed musculoskeletal dynamics with controllable representations of human-exoskeleton interaction, we can conduct fast design iterations and primary validation prior to hardware prototyping, reducing development cost and risk.

We demonstrate the use of this simulation–validation pipeline for multiple exoskeleton configurations, highlighting its ability to (i) capture the muscle-level effects of different assistance strategies, (ii) quantify their influence on gait and postural balance, and (iii) provide actionable insights for optimizing both mechanical design and control policy.

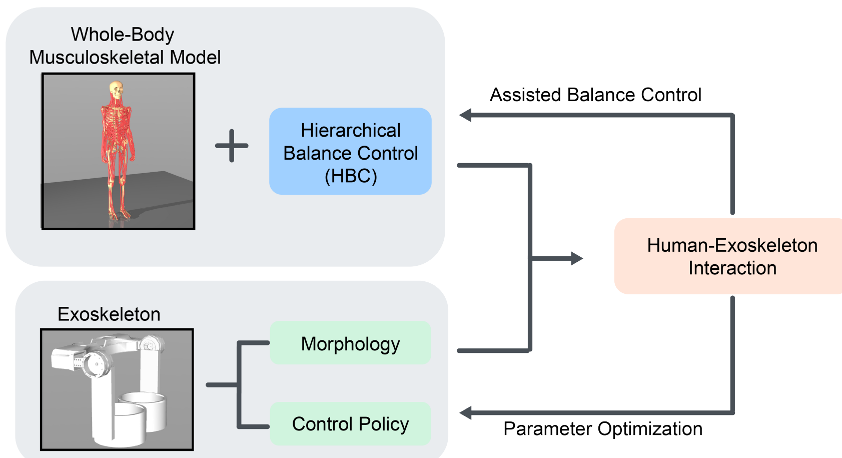

Figure 1: Control and analysis of human musculoskeletal standing and falling. HBC enables training-free balance control and efficient collection of dynamical behavior for balance analysis. Our method supports concurrent control planning for exoskeleton-assisted scenarios, facilitating integrated evaluation and optimization of human-exoskeleton interaction.

2 Related Work

2.1 Wearable robotics for Balance and Locomotion Assistance

Wearable robotics such as lower-limb exoskeletons have been widely used to improve human performance in a variety of locomotion tasks, such as walking, running, and stair climbing, for both able-bodied individuals and those with mobility impairments [13, 14, 15, 16]. They have demonstrated the ability to reduce the metabolic cost of locomotion and restore mobility for individuals with neurological disorders like stroke or spinal cord injury, aiding in gait recovery and daily activities [17, 18]. However, despite their widespread use for locomotion, their application for balance assistance has been less explored, often limited to small-scale studies in controlled lab environments [19, 20].

A major challenge in the development of exoskeletons has been the reliance on traditional control approaches that require extensive, labor-intensive human tests and the creation of handcrafted control laws for each user and activity [21, 17, 22]. This process is not only time-consuming and expensive but also creates a significant bottleneck, making it difficult to personalize controllers and scale them to diverse activities or a larger user base. This has led to a growing need for more advanced, simulation-based methods to streamline the development process [14, 16].

2.2 Musculoskeletal Simulation for Exoskeleton Development

Musculoskeletal simulations are invaluable for studying human movement biomechanics in a non-invasive way [8, 9]. They provide insights into difficult-to-measure metrics like muscle forces and joint loads, making them essential tools for designing [23], controlling [14], and evaluating assistive devices. These models have been used to estimate joint torques [24], serve as control testbeds [25], and optimize design parameters like stiffness and geometry before physical prototyping [26]. They've even been customized for specific user populations [15], offering crucial insights for creating safer and more comfortable devices [27].

Despite these advances, a critical need remains for rigorous experimental validation of model-based designs and for improving the fidelity of biomechanical models to better capture the complex dynamics of human-exoskeleton interactions [28, 27]. Many studies use reinforcement learning (RL) algorithms to optimize exoskeletons, which is time-consuming and computationally expensive, often resulting in long training times [14, 29, 16].

Our work aims to address these gaps by presenting a collaborative simulation framework utilizing a full-body human musculoskeletal model and advanced control methods [12, 30, 31] for the rapid, real-time optimization of wearable robotics configurations and control parameters across diverse bipedal balance and locomotion tasks.

3 Method

3.1 Musculoskeletal Model Dynamics

The musculoskeletal model used in this work is the MS-Human-700 model [12]. It comprises of 90 rigid body segments, 206 joints and 700 muscle-tendon units. By actuating its 700 muscle-tendon units, the model can be controlled and perform human-like tasks. The dynamics of the model can be formulated as follow:

$$M(q)\ddot{q} + c(q, \dot{q}) = J_m^T f_m(act) + J_c^T f_c + \tau_{ext}. \tag{1}$$

On the left side of the equation, q stands for generalized coordinates of joints, $M(q)$ stands for the mass distribution matrix, and $c(q, \dot{q})$ stands for Coriolis and the gravitational force. On the right side, J_m and J_c stand for Jacobian matrices that map forces to the generalized coordinates, f_c is the constraint force, $f_m(act)$ stands for actuator forces generated by muscle-tendon units determined by muscle activations (act), and τ_{ext} stands for all external torque when interacting with environments.

MS-Human-700 model is implemented in the MuJoCo physics engine [32]. The actuators of the model in this work are 700 Hill-type [33] muscles. The actuator force generated by each muscle-tendon unit, and the temporal relation between muscle activation act and the input control signal of the musculoskeletal model u can be decided by the following equations:

$$f_m(act) = f_{max} \cdot [F_l(l_m) \cdot F_v(v_m) \cdot act + F_p(l_m)]. \tag{2}$$

$$\frac{\partial act}{\partial t} = \frac{u - act}{\tau(u, act)}, \tag{3}$$

In Eq. (2), F_l and F_v represent force-length and force-velocity functions which are actuator gains, F_p is the passive force that works as actuator bias, and l_m, v_m are normalized length and normalized velocity of the muscle. f_{max} is the maximum isometric muscle force as specified in the model. In the first-order nonlinear system described by Eq. (3), muscle activation act is calculated. The time parameter τ is computed following Millard et al. [34]. τ is the time constant related to the latency in activation and deactivation.

3.2 Postural Control over Muscles

We employed the hierarchical balance control method (HBC) [30] in this work. HBC tackles the very high dimensionality of musculoskeletal model postural control by setting up high-level planning and low-level control. The high-level planning uses Model Predictive Path Integral (MPPI,

[35]) to sample low-dimensional target postures, while the low-level control maps the desired target postures back to the 700-dimensional muscle control.

3.3 Exoskeleton Control Policy and Optimization

As illustrated in Figure 1, the balance and locomotion simulation provides a testbed for strategies and designs of assistive devices, such as a hip exoskeleton, thereby enabling efficient simulation and design of exoskeletal systems. Assisted balance is simulated by applying torque actuation at the hip joints—a widely accepted approach for modeling exoskeleton effects via externally applied forces [36, 37].

We implemented a weighted postural PD control over the joint torque actuation placed at the left and right hip flexion joints. The high-level planner is adapted to plan an extra target posture indicating the overall leaning direction of the body, represented by the tilt angle of the pelvis. The control policy of the exoskeleton torque actuation is separated into two parts mixed by a weight: (1) Hip flexion joint angle PD control. (2) Postural PD control. The control is formulated as follow:

$$\tau_e^i = (1-w)\cdot(k_{p_e}\cdot(q_i^* - q_i) + k_{d_e}\cdot(0-\dot{q}_i)) + w\cdot(k_{p_t}\cdot(q_t^* - q_t) + k_{d_t}\cdot(0-\dot{q}_t)). \quad (4)$$

$i = 1, 2$ represents the left or right side respectively. τ_e^i is the torque actuation value. k_{p_e} and k_{d_e} are the joint angle PD control constants, while k_{p_t} and k_{d_t} are the postural PD control constants over the tilt angle of the pelvis. q_i^* and q_t^* are the target values of the hip joint angles and the pelvis tilt angle. w is the weight between the two PD control policies.

We found that the assistive effect of the exoskeleton was very sensitive to the k_{p_e}, k_{p_e}, k_{p_t} and w values. Therefore, we carried out Bayesian optimization (BO [38, 39]) to determine a set of parameters to ensure performance across trials. We define $\boldsymbol{x} = (k_{p_e}, k_{p_e}, k_{p_t}, w)$, and formulate the parameter search as a black-box optimization problem:

$$\max_{\boldsymbol{x}\in\mathcal{X}} f(\boldsymbol{x}) = \mathbb{E}_{\boldsymbol{x}}\left[-\sum_{t=0}^{T-1} C(s_t, u_t)\right] \quad (5)$$

where the objective $f(\boldsymbol{x})$ is the negative cumulative cost function under parameter $\boldsymbol{x}$, averaged over 5 independent trials. Given sampled data, We use Gaussian process [40] to model the posterior of the objective function, and utilize the Expected Improvement [41] acquisition function as the criterion to sample the next parameter.

4 Experiments

In the experiments, we validated the assistance provided by hip exoskeletons on the MS-Human-700 model during locomotion and balance tasks, and carried out analysis on the assistive influence of the hip exoskeleton on muscle effort.

We used MuJoCo MPC [42] in our experiments to deploy HBC in simulation time. With the very high dimensionality of the MS-Human-700 model, the simulation speed was set to 10%. A 5-second experiment will cost 50 seconds, which is much shorter than the training time of reinforcement learning methods.

4.1 Simulation of Assisted Balance

We simulated an assisted balance scenario with a hip exoskeleton, as shown in Figure 2a. We optimize the parameters of the exoskeleton control policy over 600 iterations. We employed optimized exoskeleton control parameters and validated the effectiveness of this exoskeleton control policy in a perturbation test: Models with and without the exoskeleton assistance are pushed in random directions for 3 times with intervals of 1 second. As shown in Figure 2b, balance with exoskeleton assistance achieves a higher success rate in maintaining balance under perturbation through the 5-second simulation.

In Figure 2c, we carried out an ablation study over the assisted balance by recording the muscle activations of the gluteus maximus, gluteus medius, and gluteus minimus, which play critical roles in lower-limb movement and postural stability. We observe a reduction in muscle activation levels with exoskeleton assistance, highlighting the potential of assistive devices to alleviate muscular effort and reduce metabolic cost. Such muscle-level data are difficult to obtain in traditional experimental settings due to the limitations of surface EMG and the inaccessibility of deep muscles. Our simulation pipeline enables validation of exoskeleton effects on balance maintenance. These findings suggest the utility of musculoskeletal simulations in evaluating and optimizing assistive device performance prior to costly physical prototyping and human subject testing.

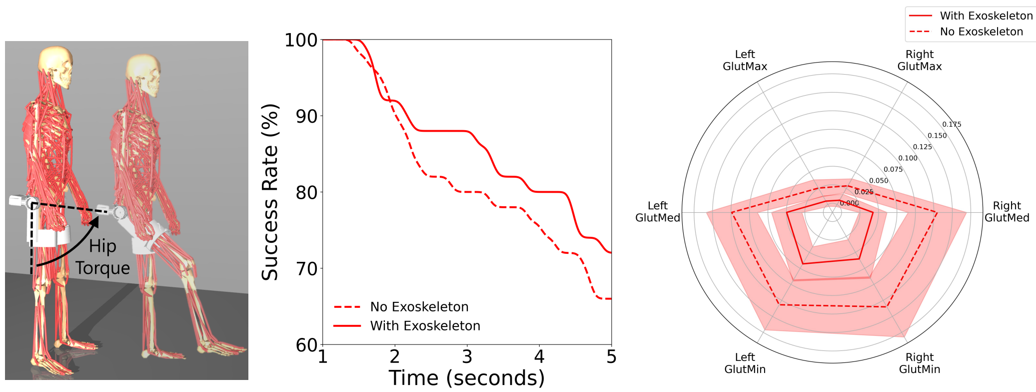

(a) (b) (c)

Figure 2: Balancing behavior under exoskeleton assistance. (a) Visualization of the hip exoskeleton device, using joint torques as actuators. (b) Success rate of maintaining standing posture, where the exoskeleton assistance helps better static balance. (c) Activation levels of gluteal muscle activations (GlutMax: gluteus maximus, GlutMed: gluteus medius, GlutMin: gluteus minimus) during the standing simulation, where muscle activation levels are reduced with exoskeleton assistance.

4.2 Simulation of Assisted Locomotion

We further designed an assisted locomotion scenario to investigate the effect of hip exoskeletons. The task cost function includes forward velocity, which encourages the model to walk forward. In the assisted locomotion task, we simulated the interaction between human body and a hip exoskeleton by a virtual representation of high-rigidity spring connections located at thighs and waist. The shape and weight distribution of the hip exoskeleton are considered in simulation, and the actuation of hip exoskeletons are modeled with actuated joints located at the motor compartment.

The control policy of the exoskeleton is crafted by hand based on proportional control and gait recognition, providing a propulsive torque to the swing leg and a counteracting torque to the stance leg during the pre-swing phase of a gait cycle. The deployment of various exoskeleton control policies, as well as testing the effects of different exoskeleton design prototypes are efficient.

5 Conclusion

In this work, we introduced a simulation-based method for the co-optimization of wearable robotics design and control with a dynamic, whole-body human musculoskeletal model. We have successfully demonstrated its capability as a validation pipeline for evaluating assistive strategies in both locomotion and balance-critical scenarios. Our experiments confirm that the method can effectively quantify the influence of different exoskeleton control policies on human muscle effort and postural stability, enabling rapid, real-time parameter tuning.

By providing a high-fidelity testbed for whole-body human-exoskeleton interaction, our work bridges a critical gap between theoretical modeling and costly physical experiments. While our

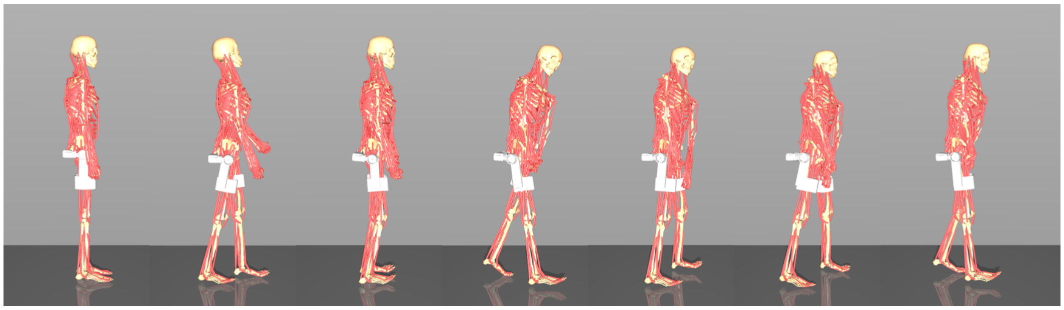

Figure 3: Human musculoskeletal model locomotion task with exoskeleton assistance. HBC enables training-free locomotion assistance demonstration and efficient parameter optimization.

current implementation utilizes a simplified representation of the exoskeleton's mechanical interface, future work will focus on incorporating more complex, realistic device models and leveraging data-driven optimization techniques to explore the co-adaptation of human and robotic controllers to accelerate the design cycle for wearable robotics. Ultimately, this simulation-driven approach holds significant promise for the development of safer, more effective, and highly personalized wearable robotics tailored to the specific needs of individual users.

References

[1] Y. Fujimoto, S. Obata, and A. Kawamura. Robust biped walking with active interaction control between foot and ground. In *Proceedings. 1998 IEEE international conference on robotics and automation (cat. no. 98ch36146)*, volume 3, pages 2030–2035. IEEE, 1998.

[2] T. Sugihara. Standing stabilizability and stepping maneuver in planar bipedalism based on the best com-zmp regulator. In *2009 IEEE International Conference on Robotics and Automation*, pages 1966–1971. IEEE, 2009.

[3] A. Takanishi, H.-o. Lim, M. Tsuda, and I. Kato. Realization of dynamic biped walking stabilized by trunk motion on a sagittally uneven surface. In *EEE International Workshop on Intelligent Robots and Systems, Towards a New Frontier of Applications*, pages 323–330. IEEE, 1990.

[4] F. Gubina, H. Hemami, and R. B. McGhee. On the dynamic stability of biped locomotion. *IEEE Transactions on Biomedical Engineering*, (2):102–108, 1974.

[5] J. Pratt and G. Pratt. Intuitive control of a planar bipedal walking robot. In *Proceedings. 1998 IEEE international conference on robotics and automation (Cat. No. 98CH36146)*, volume 3, pages 2014–2021. IEEE, 1998.

[6] R. Katoh and M. Mori. Control method of biped locomotion giving asymptotic stability of trajectory. *Automatica*, 20(4):405–414, 1984.

[7] F. J. Valero-Cuevas, H. Hoffmann, M. U. Kurse, J. J. Kutch, and E. A. Theodorou. Computational models for neuromuscular function. *IEEE reviews in biomedical engineering*, 2: 110–135, 2009.

[8] 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*, 54(11):1940–1950, 2007.

[9] C. Vittorio, W. Huawei, D. Guillaume, S. Massimo, and K. Vikash. Myosuite – a contact-rich simulation suite for musculoskeletal motor control. `https://github.com/myohub/myosuite`, 2022. URL `https://arxiv.org/abs/2205.13600`.

[10] Y. Sui, K. ho Kim, and J. W. Burdick. Quantifying performance of bipedal standing with multichannel emg. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3891–3896. IEEE, 2017.

[11] K. Akbaş, C. Mummolo, and X. Zhou. Characterization of human balance through a reinforcement learning-based muscle controller. *arXiv preprint arXiv:2308.04462*, 2023.

[12] C. Zuo, K. He, J. Shao, and Y. Sui. Self model for embodied intelligence: Modeling full-body human musculoskeletal system and locomotion control with hierarchical low-dimensional representation. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 13062–13069, 2024. doi:10.1109/ICRA57147.2024.10610081.

[13] D. Pinto-Fernandez, D. Torricelli, M. del Carmen Sanchez-Villamanan, F. Aller, K. Mombaur, R. Conti, N. Vitiello, J. C. Moreno, and J. L. Pons. Performance evaluation of lower limb exoskeletons: a systematic review. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 28(7):1573–1583, 2020.

[14] S. Luo, M. Jiang, S. Zhang, J. Zhu, S. Yu, I. Dominguez Silva, T. Wang, E. Rouse, B. Zhou, H. Yuk, et al. Experiment-free exoskeleton assistance via learning in simulation. *Nature*, 630 (8016):353–359, 2024.

[15] Y. Lu, Y. Huang, R. Yang, Y. Wang, Y. Ikegami, Y. Nakamura, and Q. Wang. A human-prosthesis coupled musculoskeletal model for transtibial amputees. *IEEE Transactions on Biomedical Engineering*, 2025.

[16] M. Sharifi, S. Tripathi, Y. Chen, Q. Zhang, and M. Tavakoli. Reinforcement learning methods for assistive and rehabilitation robotic systems: A survey. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 2025.

[17] P. Slade, M. J. Kochenderfer, S. L. Delp, and S. H. Collins. Personalizing exoskeleton assistance while walking in the real world. *Nature*, 610(7931):277–282, 2022.

[18] N. Hankov, M. Caban, R. Demesmaeker, M. Roulet, S. Komi, M. Xiloyannis, A. Gehrig, C. Varescon, M. R. Spiess, S. Maggioni, et al. Augmenting rehabilitation robotics with spinal cord neuromodulation: A proof of concept. *Science robotics*, 10(100):eadn5564, 2025.

[19] V. Monaco, P. Tropea, F. Aprigliano, D. Martelli, A. Parri, M. Cortese, R. Molino-Lova, N. Vitiello, and S. Micera. An ecologically-controlled exoskeleton can improve balance recovery after slippage. *Scientific reports*, 7(1):46721, 2017.

[20] A. Emmens, E. Van Asseldonk, M. Masciullo, M. Arquilla, I. Pisotta, N. L. Tagliamonte, F. Tamburella, M. Molinari, and H. Van Der Kooij. Improving the standing balance of paraplegics through the use of a wearable exoskeleton. In *2018 7th IEEE International Conference on Biomedical Robotics and Biomechatronics (Biorob)*, pages 707–712. IEEE, 2018.

[21] T. Yan, M. Cempini, C. M. Oddo, and N. Vitiello. Review of assistive strategies in powered lower-limb orthoses and exoskeletons. *Robotics and Autonomous Systems*, 64:120–136, 2015.

[22] C. Siviy, L. M. Baker, B. T. Quinlivan, F. Porciuncula, K. Swaminathan, L. N. Awad, and C. J. Walsh. Opportunities and challenges in the development of exoskeletons for locomotor assistance. *Nature biomedical engineering*, 7(4):456–472, 2023.

[23] D. D. Molinaro, K. L. Scherpereel, E. B. Schonhaut, G. Evangelopoulos, M. K. Shepherd, and A. J. Young. Task-agnostic exoskeleton control via biological joint moment estimation. *Nature*, 635(8038):337–344, 2024.

[24] S. K. Lam and I. Vujaklija. Joint torque prediction via hybrid neuromusculoskeletal modelling during gait using statistical ground reaction estimates: an exploratory study. *Sensors*, 21(19): 6597, 2021.

[25] F. Mo, Q. Zhang, H. Zhang, J. Long, Y. Wang, G. Chen, and J. Ye. A simulation-based framework with a proprioceptive musculoskeletal model for evaluating the rehabilitation exoskeleton system. *Computer Methods and Programs in Biomedicine*, 208:106270, 2021.

[26] S. Yu, L. Liu, S. Zhang, A. Di Lallo, J. Zhu, Q. Wu, G. Zuo, X. Zhou, and H. Su. Controlling negative and positive power for efficiency enhancement and muscle strain mitigation during squatting with a portable knee exoskeleton. *Annals of Biomedical Engineering*, 53(6):1344–1358, 2025.

[27] V. Firouzi, A. Seyfarth, S. Song, O. von Stryk, and M. Ahmad Sharbafi. Biomechanical models in the lower-limb exoskeletons development: A review. *Journal of NeuroEngineering and Rehabilitation*, 22(1):12, 2025.

[28] D. Scherb, S. Wartzack, and J. Miehling. Modelling the interaction between wearable assistive devices and digital human models—a systematic review. *Frontiers in bioengineering and biotechnology*, 10:1044275, 2023.

[29] M. W. Sang, J. Narayan, B. Omarali, and A. A. Faisal. Towards safer rehabilitation: Improving gait trajectory tracking for lower limb exoskeletons using offline reinforcement learning. In *2025 International Conference On Rehabilitation Robotics (ICORR)*, pages 577–582. IEEE, 2025.

[30] C. Ma, Y. Wei, C. Zuo, C. Zhang, and Y. Sui. Bipedal balance control with whole-body musculoskeletal standing and falling simulations. *arXiv preprint arXiv:2506.09383*, 2025.

[31] Y. Wei, S. Zhuang, V. Zhuang, and Y. Sui. Motion control of high-dimensional musculoskeletal systems with hierarchical model-based planning. *arXiv preprint arXiv:2505.08238*, 2025.

[32] E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033, 2012. doi:10.1109/IROS.2012.6386109.

[33] F. E. Zajac. Muscle and tendon: properties, models, scaling, and application to biomechanics and motor control. *Critical reviews in biomedical engineering*, 17(4):359–411, 1989.

[34] M. Millard, T. Uchida, A. Seth, and S. L. Delp. Flexing computational muscle: modeling and simulation of musculotendon dynamics. *Journal of biomechanical engineering*, 135(2): 021005, 2013.

[35] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou. Aggressive driving with model predictive path integral control. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1433–1440, 2016. doi:10.1109/ICRA.2016.7487277.

[36] T. K. Uchida, A. Seth, S. Pouya, C. L. Dembia, J. L. Hicks, and S. L. Delp. Simulating ideal assistive devices to reduce the metabolic cost of running. *PloS one*, 11(9):e0163417, 2016.

[37] C. L. Dembia, A. Silder, T. K. Uchida, J. L. Hicks, and S. L. Delp. Simulating ideal assistive devices to reduce the metabolic cost of walking with heavy loads. *PloS one*, 12(7):e0180320, 2017.

[38] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2015.

[39] P. I. Frazier. A tutorial on bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018.

[40] C. E. Rasmussen. Gaussian processes in machine learning. In *Summer school on machine learning*, pages 63–71. Springer, 2003.

[41] D. R. Jones, M. Schonlau, and W. J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13:455–492, 1998.

[42] T. Howell, N. Gileadi, S. Tunyasuvunakool, K. Zakka, T. Erez, and Y. Tassa. Predictive sampling: Real-time behaviour synthesis with mujoco, 2022. URL `https://arxiv.org/abs/2212.00541`.