MuscleMimic: Unlocking full-body musculoskeletal motor learning at scale

Chengkun Li[1,†] Cheryl Wang[2,†] Bianca Ziliotto[1] Merkourios Simos[1]
Guillaume Durandau[2] Alexander Mathis[1,*]
[1]EPFL, Switzerland [2]McGill University, Canada
[*]alexander.mathis@epfl.ch [†]Equal contribution

Abstract

The computational cost of biomechanically accurate simulation and the scarcity of validated, open full-body models hinders learning policies for muscle-driven musculoskeletal models. We present MuscleMimic, an open-source framework for scalable motion imitation learning with physiologically realistic, muscle-actuated humanoids. MuscleMimic provides two validated musculoskeletal models, a fixed-root upper-body model (126 muscles) for bimanual manipulation and a full-body model (416 muscles) for locomotion, together with a retargeting pipeline that maps SMPL-format motion capture data onto musculoskeletal structures while preserving kinematic and dynamic consistency. Leveraging massively parallel GPU simulation, the framework achieves one order-of-magnitude training speedups over prior CPU-based approaches, enabling a single generalist policy to be trained on hundreds of diverse motions within days. Biomechanical validation against experimental data demonstrates strong kinematic agreement for walking (r=0.90) and running (r=0.79), while muscle activation analysis reveals both the promise and fundamental challenges of achieving physiological fidelity through kinematic imitation. Code, musculoskeletal models, policy checkpoints, and retargeted datasets are available at https://github.com/amathislab/musclemimic.

1. Introduction

Perceiving and understanding human motion from visual data is a cornerstone of computer vision. The past decade has seen tremendous progress in 3D human pose estimation, transitioning from sparse 2D eypoints to dense 3D surface meshes driven by parametric models like SMPL [14]. This progress has yielded large-scale kinematic datasets such as AMASS [16], which capture a vast repertoire of human behaviors. However, visual kinematics alone are insufficient for a physical understanding of movement: they ignore mass, inertia, ground reaction forces, and internal actuation, frequently resulting in physically implausible artifacts such as foot sliding, ground penetration, and dynamically impossible accelerations.

To bridge the gap between visual perception and physical embodiment, many recent works leverage reinforcement learning to imitate visual reference motions [15, 17, 24]. Yet these physics-based approaches typically rely on simplified torque-driven models, circumventing the human neuromotor system. If Embodied AI is to move beyond superficial motion synthesis toward a true physical understanding, we must transition to anatomically grounded musculoskeletal (MSK) models [3, 11]. However, simulating complex muscle pathways and collision dynamics on traditional CPU-bound engines is exceptionally slow, confining MSK motor learning to narrow movement repertoires [5, 6, 8, 13, 19, 21].

We present **MuscleMimic**, an open-source framework that bridges the scale of computer vision motion data with the fidelity of biomechanical simulation. Our contributions are: **(1)** Two highly detailed MSK models (hereafter interchangably referred to as embodiments): a bimanual upper-body model (126 muscles) and a full-body model (416 muscles); **(2)** A retargeting pipeline that maps standard SMPL-format motion capture data, the dominant representation in computer vision and computer graphics, into physically viable MSK trajectories, enabling direct use of existing datasets like AMASS; and **(3)** A massively parallel training environment that achieves 1.4×10^4 steps/sec on a single GPU, fueling a single generalist policy to be trained on hundreds of diverse motions within hours rather than days.

2. Method

2.1. Musculoskeletal Models

MuscleMimic provides two complementary musculoskeletal embodiments (Table 1), both built upon established MyoSuite components [4]. Muscle actuators follow the Hill-type model [10] with first-order nonlinear activation dynamics, where control signals are transformed through state-dependent time constants (τ_{act}=0.01s, τ_{deact}=0.04s). Both models enforce bilateral symmetry in joint constraints, muscle moment arms, and force–length curves.

BimanualMuscle is designed for upper-body manipula-

Table 1. Overview of the two musculoskeletal embodiments. *: finger muscles disabled.

Model	Type	Joints	Muscles	DoFs	Focus
BimanualMuscle	Fixed-base	76 (36*)	126 (64*)	54 (14*)	Upper body
MyoFullBody	Free-root	123 (83*)	416 (354*)	72 (32*)	Full body

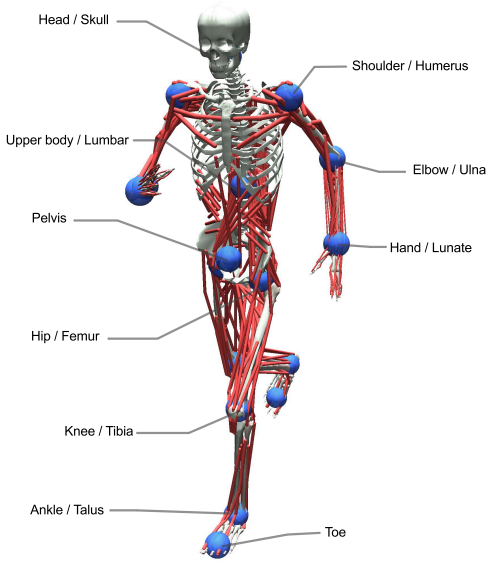

Figure 1. The 17 mimic sites used for full-body motion imitation, with corresponding anatomical landmarks. For BimanualMuscle, a subset of 7 upper-limb sites is used.

tion with a fixed thorax, 126 Hill-type muscle actuators with bilateral symmetry, and 7 mimic sites for motion tracking. **MyoFullBody** provides a complete full-body system with 123 joints from pelvis to fingertips, 416 muscles, 17 mimic sites at key anatomical landmarks (Fig. 1), and comprehensive collision detection supporting both self-collision and environment contact.

2.2. SMPL-to-MSK Retargeting Pipeline

A key contribution is our retargeting pipeline that bridges SMPL-format motion data, the standard in computer vision, with biomechanically constrained musculoskeletal simulation, enabling large-scale training on existing motion datasets like AMASS [16].

The pipeline (Fig. 2) consists of three stages. **Preprocessing:** AMASS motion sequences, parameterized by SMPL shape ($\boldsymbol{\beta}$) and pose parameters, undergo shape fitting in a T-pose to estimate body proportions, a global scale s, and joint-wise offsets $(\Delta\mathbf{p}, \Delta\mathbf{R})$ that map SMPL morphology to the MSK model. **Inverse kinematics:** Scaled motions are retargeted via two complementary methods: *Mocap-Body*, using MuJoCo's internal IK, and *GMR-Fit* [2], which enforces model-defined joint constraints and reduces posture jumps between frames. **Post-processing:** Retargeted trajectories are interpolated to 100 Hz control frequency and corrected for ground penetration and floating artifacts.

GMR-Fit achieves substantially lower joint-limit violations (0.27% vs. 12.26%) and tendon jump rates (3.2% vs. 30.1%) compared to Mocap-Body on the KINESIS dataset [20], while Mocap-Body is $\sim 3\times$ faster. In practice, we use GMR-Fit due to its superior kinematic fidelity.

3. Results

3.1. Training Efficiency

Training muscle-actuated MSK models has traditionally been hampered by simulation cost. Our pipeline uses MuJoCo Warp [7] for massively parallel GPU simulation with comprehensive contact handling, coupled with a JAX-based training framework extending LocoMuJoCo [1] for end-to-end learning on a single GPU.

We benchmark end-to-end training throughput for MyoFullBody with a 100M-parameter policy on a single NVIDIA H100 80GB GPU (Fig. 3). We observe that training steps per second (SPS) increase from 174 SPS at n=16 to 1.4×10^4 SPS at n=8,192 parallel environments. One billion environment steps complete in $\sim$20 hours on a single GPU. For reference, KINESIS [20] required $\sim$10 days on 128 CPU cores with an A100 GPU for a 290-actuator model; while the embodiments and hardware differ, our 416-actuator MyoFullBody model trained to convergence in under 3 days on a single GPU.

3.2. Motion Imitation

Qualitative results. By imitating human motion capture data, our generalist policy acquired diverse motor skills while controlling over 400 independent muscle actuators. The MyoFullBody model produced diverse locomotion skills including walking, running, turning, and dancing (Fig. 4). The BimanualMuscle model reproduced a broad range of upper-body movements spanning sports, object interactions, and daily activities. Beyond the generalist policy, MuscleMimic supports fine-tuning: gentle motions such as dancing require $<$100M steps, while complex dynamics like vertical jumping and kick-twists require $\sim$1B steps.

Quantitative results. We evaluated both embodiments on training and test sets using pretrained checkpoints with three seeds (Table 2). A motion is considered successful if the agent completes it without early termination, which is triggered when the mean mimic site deviation relative to the root exceeds 0.5 m (MyoFullBody, 17 sites) or 0.25 m (BimanualMuscle, 6 sites). MyoFullBody achieved $95.5\pm0.4\%$ success rate on 972 training motions and $92.6\pm0.0\%$ on 108 held-out test motions, with a mean joint angle error of 6.63° and relative site position error of 2.26 cm. BimanualMuscle achieved even higher performance ($>$98% success) on its 1770-motion training set.

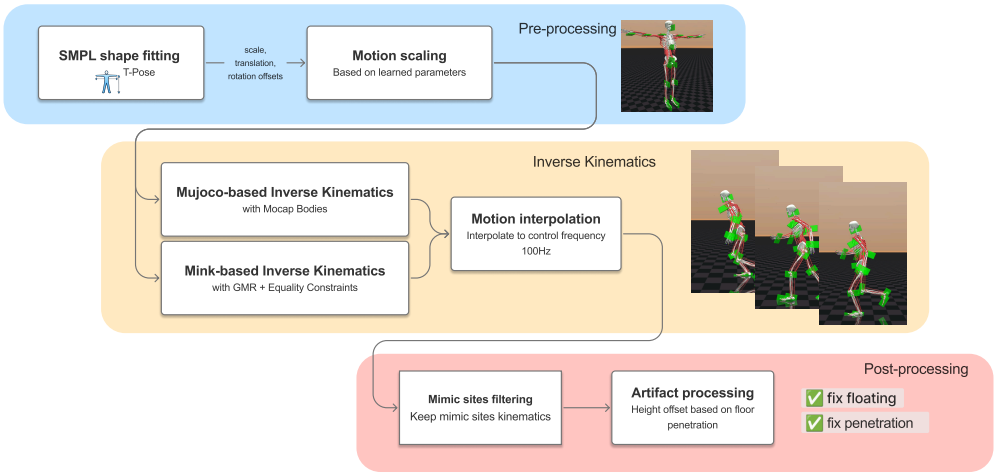

Figure 2. Overview of the SMPL-to-MSK retargeting pipeline. Human motion in SMPL format is shape-fitted, retargeted via inverse kinematics (Mocap-Body or GMR-Fit), and post-processed to correct artifacts, producing biomechanically valid reference trajectories.

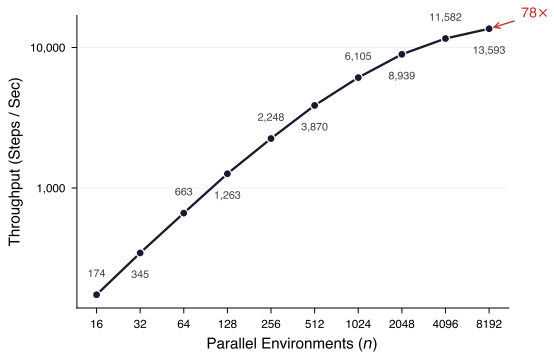

Figure 3. End-to-end training throughput (steps/sec) for MyoFullBody with a 100M-parameter policy as the number of parallel environments scales from 16 to 8192 on a single H100 GPU.

Table 2. Key validation metrics for both embodiments (GMR-Fit retargeting, N=3 seeds). Full results include 972/108 training/test motions for MyoFullBody and 1770/312 for BimanualMuscle.

Metric	MyoFullBody		BimanualMuscle	
	Train	Test	Train	Test
Success rate (%) ↑	95.5±0.4	92.6±0.0	98.6±0.0	99.5±0.0
Joint angle err (°) ↓	6.67±0.00	6.63±0.01	6.19±0.01	6.09±0.02
Rel. site err (cm) ↓	2.26±0.00	2.28±0.02	1.90±0.01	1.84±0.01

4. Biomechanical Validation

To verify biomechanical fidelity beyond training metrics, we evaluated against independent experimental data on joint kinematics, ground reaction forces (GRF), and electromyography (EMG) [9, 18] for two datasets [12, 23]. Our validation focused on the canonical features of the gait cycle [22] rather than trial-specific trajectory tracking to ensure our policy captures the fundamental physiological principles of human locomotion.

Kinematics and kinetics. We evaluated the pretrained MyoFullBody policy on five AMASS walking sequences, comparing against treadmill walking at 1.2 m/s [23] and level walking [12], both averaged across nine participants. Signals were aligned by GRF onset and normalized to a single gait cycle. As shown in Fig. 5, simulated joint kinematics achieve mean correlations of r=0.9 (treadmill & level walking) with experimental data, demonstrating accurate reproduction of the stereotyped hip–knee–ankle coordination pattern and the characteristic double-peaked GRF profile. For running at 1.8 m/s, the model achieved r=0.79, capturing the key differences from walking, including the single-peaked GRF and pronounced swing-phase knee flexion.

Muscle activation analysis. We compared synthetic muscle activations with EMG recordings during walking [12, 23] (Fig. 6). A central finding was that strong kinematic imitation does not automatically guarantee physiologically faithful muscle activations. Across independently trained policies, the observed muscle–EMG correlations span 0.2 to 0.6; this range reflects variability across muscles and across training runs with different random seeds and different reward trade-offs between kinematic tracking and energy regularization. Under a similar motion imitation formulation, KINESIS reported correlations of approximately 0 to 0.45 and shows that non-imitation baselines yield substantially weaker EMG alignment than imitation-based controllers [20], while our kinematic tracking error is lower. Repeated experiments reliably produce positive correlations, but the specific magnitude varied substantially across policies and muscles. This

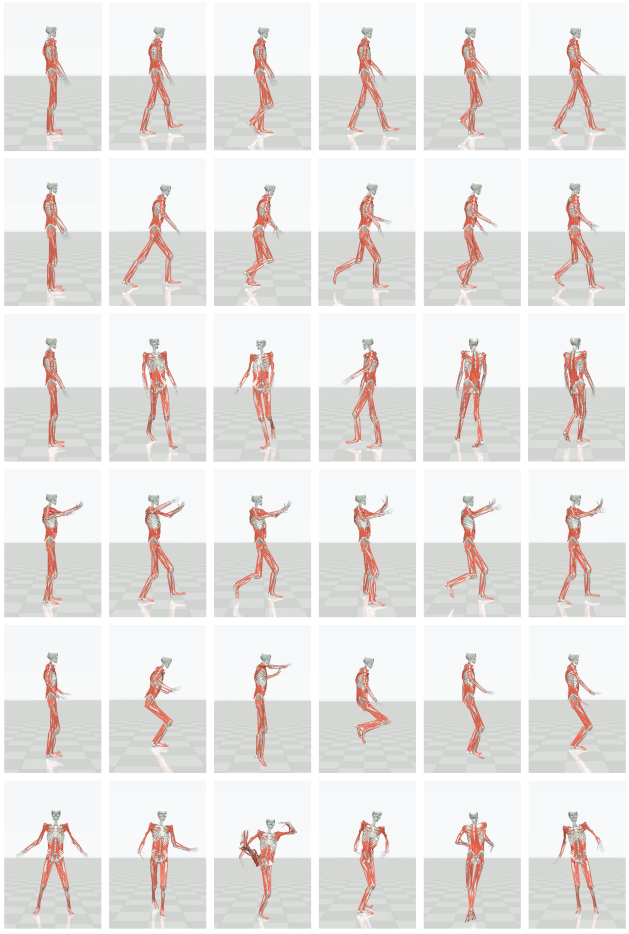

Figure 4. Motion snapshots from trained MyoFullBody policies. Top to bottom: walking, running, turning, dancing, vertical jumping, and kick with 360° twist.

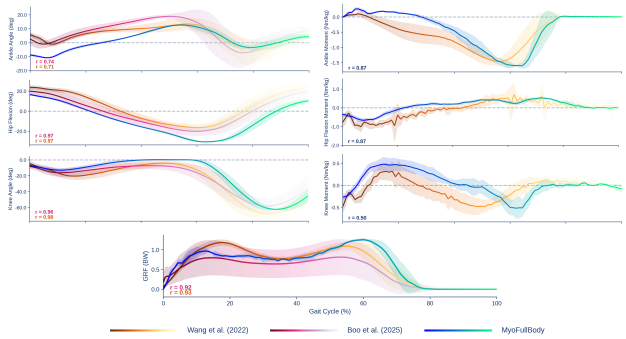

Figure 5. Left lower-limb joint kinematics over a walking gait cycle, comparing experimental human data (treadmill [23] and level walking [12]) with MyoFullBody simulation. Mean correlations: r=0.9 (treadmill & level walking). Shaded areas represent ± 1 standard deviation across all trials.

variability is a principled consequence of muscle redundancy: many distinct coordination strategies can produce mechanically equivalent joint kinematics, and a controller optimized for kinematic accuracy has no incentive to converge on the particular strategy employed by human subjects.

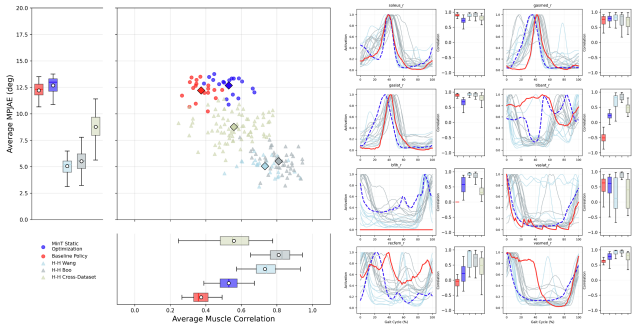

Figure 6. Physiological plausibility of synthetic muscle activations during walking. Left: average muscle–EMG correlation vs. lower-limb MPJAE; triangles denote human-to-human pairs, dots denote model-to-human pairs. Right: gait-cycle-averaged activation patterns and per-muscle correlation values.

5. Discussion and Conclusion

We introduced MuscleMimic, which enables scalable motion imitation learning with two novel physiologically realistic musculoskeletal models. By leveraging GPU-accelerated simulation with massive parallelism, we achieve improvements in training speed while maintaining comprehensive collision handling.

While our framework demonstrates promising alignment with experimental kinematic data, musculoskeletal models remain approximations of biological reality. The Hill-type muscle model simplifies tendon elasticity and fiber recruitment. Muscle redundancy remains a fundamental challenge: kinematic accuracy alone does not ensure physiologically faithful activations. SMPL-based retargeting assumes a generic morphology, potentially introducing systematic biases for atypical anthropometrics. Our biomechanical validation currently focuses on walking and running; extending to the full range of demonstrated motions (dancing, jumping, kick-twists) requires corresponding experimental datasets.

By open-sourcing all code, models, checkpoints, and the retargeted AMASS dataset, we hope to enable the research community to build on these models and to foster collaboration at the intersection of computer vision, biomechanics and neuroscience.

References

[1] Firas Al-Hafez, Guoping Zhao, Jan Peters, and Davide Tateo. Locomujoco: A comprehensive imitation learning benchmark for locomotion. *arXiv preprint arXiv: 2311.02496*, 2023. 2

[2] Joao Pedro Araujo, Yanjie Ze, Pei Xu, Jiajun Wu, and C. Karen Liu. Retargeting matters: General motion retargeting for humanoid motion tracking. *arXiv preprint arXiv:2510.02252*, 2025. 2

[3] Andy Bonnetto, Sepideh Mamooler, Chengkun Li, and Alexander Mathis. The behavior biopsy: Interpreting animal behavior as embodied, situated, and hierarchical. *Current Opinion in Neurobiology*, 98:103201, 2026. 1

[4] Vittorio Caggiano, Huawei Wang, Guillaume Durandau, Massimo Sartori, and Vikash Kumar. Myosuite–a contact-rich simulation suite for musculoskeletal motor control. *arXiv preprint arXiv:2205.13600*, 2022. 1

[5] Alberto Silvio Chiappa, Pablo Tano, Nisheet Patel, Abigail Ingster, Alexandre Pouget, and Alexander Mathis. Acquiring musculoskeletal skills with curriculum-based reinforcement learning. *Neuron*, 112(23):3969–3983, 2024. 1

[6] Menthy Denayer, Eligia Alfio, María Alejandra Díaz, Massimo Sartori, Friedl De Groote, Kevin De Pauw, and Tom Verstraten. A prisma systematic review through time on predictive musculoskeletal simulations. *Journal of Neuro-Engineering and Rehabilitation*, 22(1), 2025. 1

[7] Google DeepMind. Mujoco warp: Gpu-optimized version of the mujoco physics simulator. https://github.com/google-deepmind/mujoco_warp, 2024. 2

[8] Kaibo He et al. Dynsyn: Dynamical synergistic representation for efficient learning and control in overactuated embodied systems. *arXiv preprint arXiv:2407.11472*, 2024. 1

[9] Jennifer L. Hicks, Thomas K. Uchida, Ajay Seth, Apoorva Rajagopal, and Scott L. Delp. Is my model good enough? best practices for verification and validation of musculoskeletal models and simulations of movement. *Journal of Biomechanical Engineering*, 137(2), 2015. 3

[10] A. V. Hill. The heat of shortening and the dynamic constants of muscle. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, 126(843):136–195, 1938. 1

[11] Marilyn Keller, Keenon Werling, Soyong Shin, Scott Delp, Sergi Pujades, C. Karen Liu, and Michael J. Black. From skin to skeleton: Towards biomechanically accurate 3d digital humans. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)*, 42(6):253:1–253:15, 2023. 1

[12] Seungbum Koo, Junyo Boo, Dongwook Seo, and Minseung Kim. Comprehensive human locomotion and electromyography dataset: Gait120, 2025. 3, 4

[13] Seunghwan Lee, Moonseok Park, Kyoungmin Lee, and Jehee Lee. Scalable muscle-actuated human simulation and control. *ACM Transactions on Graphics*, 38(4):1–13, 2019. 1

[14] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. *SMPL: A Skinned Multi-Person Linear Model*. Association for Computing Machinery, New York, NY, USA, 1 edition, 2023. 1

[15] Zhengyi Luo, Jinkun Cao, Kris Kitani, Weipeng Xu, et al. Perpetual humanoid control for real-time simulated avatars. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10895–10904, 2023. 1

[16] Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. Amass: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5442–5451, 2019. 1, 2

[17] Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel Van de Panne. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions On Graphics (TOG)*, 37(4):1–14, 2018. 1

[18] Apoorva Rajagopal, Christopher L. Dembia, Matthew S. DeMers, Denny D. Delp, Jennifer L. Hicks, and Scott L. Delp. Full-body musculoskeletal model for muscle-driven simulation of human gait. *IEEE Transactions on Biomedical Engineering*, 63(10):2068–2079, 2016. 3

[19] Pierre Schumacher, Daniel F.B. Haeufle, Dieter Büchler, Syn Schmitt, and Georg Martius. Dep-rl: Embodied exploration for reinforcement learning in overactuated and musculoskeletal systems. *International Conference on Learning Representations (ICLR)*, 2023. 1

[20] Merkourios Simos, Alberto Silvio Chiappa, and Alexander Mathis. Reinforcement learning-based motion imitation for physiologically plausible musculoskeletal motor control. *arXiv preprint arXiv:2503.14637*, 2025. 2, 3

[21] Nikolaos Smyrnakis, Tasos Karakostas, and R. James Cotton. Advancing monocular video-based gait analysis using motion imitation with physics-based simulation. In *2024 10th IEEE RAS/EMBS International Conference for Biomedical Robotics and Biomechatronics (BioRob)*, pages 102–108, 2024. 1

[22] Thomas K. Uchida and Scott L. Delp. *Biomechanics of Movement: The Science of Sports, Robotics, and Rehabilitation*. MIT Press, Cambridge, MA, 2020. 3

[23] Huawei Wang, Akash Basu, Guillaume Durandau, and Massimo Sartori. Comprehensive kinetic and emg dataset of daily locomotion with 6 types of sensors, 2022. 3, 4

[24] Jungdam Won, Deepak Gopinath, and Jessica Hodgins. A scalable approach to control diverse behaviors for physically simulated characters. *ACM Transactions on Graphics*, 39(4), 2020. 1