VolMusCo: Learning Full-Body Motion Control with Volumetric Muscles

Wangyingqi Ding[1*] Yusen Feng[2*] Kaibo He[3] Bingchen Han[2] Libin Liu[2]
[1]Tsinghua University [2]Peking University [3]Clone Incorporated
[*]Equal contribution

Abstract

We present a simulation and learning framework for full-body motion control driven by volumetric muscles. Our approach establishes two-way coupling between muscles and the skeleton: muscle activations induce contraction along embedded fibers, deforming the volumetric tissue and generating forces that are transmitted to the skeleton through muscle attachments and wrapping contacts. In turn, skeletal motion changes the attachment and wrapping geometry, influencing tissue deformation and force transmission. This coupling brings volumetric deformation directly into the controlled dynamics. To make learning computationally efficient, we introduce a simplified tetrahedral muscle representation that preserves essential volumetric deformation and force transmission while substantially reducing simulation cost. Building on this representation, we develop a model-based reinforcement learning framework that learns a differentiable world model from physical transitions and uses predicted future motion to improve the control policy. On LaFAN1 motion tracking, our method achieves a tracking success rate of 97.9% and a mean local-body position error of 5.52 cm within 10 h of training. Project page and video demonstrations are available at https://inckieultra.github.io/VolMusCo-Project-Page/.

1. Introduction

Human motion emerges from the coordinated interaction between skeletal dynamics and muscle contraction. Muscle deformation and skeletal motion are mutually dependent and evolve through a coupled physical process. Neural activations induce contraction along muscle fibers, deforming the surrounding tissue and generating forces that drive the skeleton, while also shaping the body's appearance. Conversely, skeletal motion changes muscle attachment and wrapping configurations, which in turn influence tissue deformation and force transmission.

Existing muscle-driven character control predominantly relies on Hill-type muscle models, which represent the character's musculature as hundreds of one-dimensional, massless actuators attached to rigid skeletal segments [18, 33, 60]. Recent methods have advanced motion control with these models through hierarchical coordination [21], muscle-feasible motion tracking [37], and GPU-parallel imitation learning [25]. However, this simplified representation cannot explicitly capture volumetric tissue deformation, such as transverse bulging, or interactions between muscles and surrounding anatomical structures. Accounting for dynamic effects of muscle motion, such as inertia, also requires additional modeling [13, 53].

Volumetric muscle models provide the spatial degrees of freedom needed to capture these deformations and interactions. They have been widely used in computer graphics to simulate detailed muscle deformation and anatomical interactions [22, 32, 49]. In many such applications, muscle deformation is computed in response to prescribed skeletal motion, forming a one-way coupling from the skeleton to the muscle tissue. Recent work has explored learning motion control using one-dimensional muscle models and applying the resulting muscle activations to volumetric simulations [26]. However, the coupling remains one-way, as volumetric simulation is outside the control loop and tissue responses do not affect skeletal motion during control. Prior work has also established two-way coupling between volumetric muscles and skeletal dynamics [6, 20], allowing muscle deformation and its mechanical reactions to jointly determine skeletal motion. However, the high dimensionality of the coupled system makes learning-based control challenging, and existing approaches have therefore focused on more specialized settings, such as upper-body manipulation or task-specific trajectory optimization, rather than scalable full-body motion control.

In this paper, we present VolMusCo, a simulation and learning framework for full-body motion control driven by volumetric muscles. Our approach establishes two-way coupling between muscle tissue and skeletal dynamics, integrating tissue deformation and its mechanical effects directly into the control loop.

A major challenge in learning under such coupled dy-

namics is the computational cost of repeatedly solving volumetric tissue equilibrium during reinforcement learning. To make physical simulation efficient, we introduce a simplified tetrahedral muscle representation that preserves essential volumetric deformation and fiber-based actuation while substantially reducing the cost of tissue simulation. We also adopt a model-based reinforcement learning framework that learns a differentiable world model from physical transitions collected in the volumetric simulator. The learned model predicts future physical states under muscle excitations, allowing the policy and critic to improve tracking behavior through short predicted trajectories without repeatedly executing expensive tissue solves.

We evaluate VolMusCo on full-body motion tracking using LaFAN1 dataset [15] and additional fitness and object-manipulation motions. On LaFAN1 clips, our method achieves a tracking success rate of 97.9% and a mean local-body position error of 5.52 cm within 10 hours of training. We also conduct controlled experiments to evaluate the fidelity of the simplified muscle representation, the effect of deformation-dependent force feedback, and the learning efficiency gained through world-model prediction. The paper makes two contributions:

- We present a new learning framework for full-body motion control driven by two-way coupled volumetric muscles. To our knowledge, this is the first framework to enable such control.
- We introduce a computationally efficient design that combines a simplified tetrahedral muscle representation with model-based reinforcement learning, enabling scalable learning across a diverse range of full-body motions.

2. Related Work

Learning muscle-driven control. Early controllers combine muscle reflexes with legged mechanics [11], optimize biologically motivated locomotion objectives [51], and adapt muscle control across biped morphologies [9]. Many-muscle locomotion [23] extends control to detailed humanoid musculature. OpenSim [4] and detailed musculoskeletal models [39] support internal-force analysis. OpenSim Moco [5] enables muscle-driven motion through trajectory optimization. MASS [21] instead learns skeletal coordination and muscle recruitment at separate levels. Generative GaitNet [35] conditions control on anatomy and gait requirements; Bidirectional GaitNet [36] also estimates anatomical conditions from gait. MAGNET [37] addresses diverse-motion tracking through muscle-feasible torque clipping, a muscle coordination network, and staged adaptation. Other approaches address overactuation through structured exploration, as in DEP-RL [45], or state-dependent synergistic action representations, as in DynSyn [16]. MyoSuite [2] supplies contact-rich musculoskeletal tasks. KINESIS [46] learns muscle-driven locomotion through imitation, while MuscleMimic [25] combines full-body assets, retargeting, and GPU-parallel imitation learning. MS-Emulator [54] learns muscle-driven imitation with adversarial rewards and a value-guided flow policy. These systems broaden motion coverage, but kinematics alone leave muscle recruitment ambiguous. Muscles in Action [3] and Muscles in Time [42] study muscle–motion relationships through measured or simulated labels. MUSIC [56] distills high-frequency tracking for dexterous piano performance into a low-dimensional controller reusable by a slower music-conditioned policy. Wang et al. [52] train direct excitation policies for athletic locomotion using task rewards and high-throughput simulation, without motion demonstrations. MuscleVAE [7] is the closest model-based muscle-control precedent, incorporating muscle-length control, fatigue, and a generative skill representation. We reuse the model-based control principle without its generative VAE; our focus is deformation-dependent force feedback to the skeleton.

World models for motion tracking. DeepMimic [38] establishes reference-guided reinforcement learning for simulated skills. PULSE [29] learns reusable humanoid motion representations, and MaskedMimic [50] formulates control from partial motion descriptions as motion inpainting. BeyondMimic [28] combines agile tracking with guided-diffusion control; SONIC [30] scales a generalist tracker through model capacity, motion data, and compute. Learned dynamics provide another route to reusing experience. Dyna [48] integrates environment interaction and model-based updates, while model-based policy optimization [19] uses short model rollouts branched from real data to limit model bias. Dreamer [14] learns behavior using future states predicted in latent space. For physical animation, SuperTrack [8] optimizes tracking through learned transition derivatives, ControlVAE [58] jointly learns dynamics and generative control, and MoConVQ [59] learns scalable discrete control representations. SuperMimic [27] trains a shared world model and multi-motion actor together, with a critic extending short model-predicted trajectories. MorphMimic [57] studies sharing across humanoid morphologies. We follow the same model-based learning paradigm and extend it to motion control with volumetric muscles. SHAC [55] and AHAC [10] additionally motivate our attention to finite-horizon gradients and stiff contact dynamics.

Volumetric muscle simulation and control. Teran et al. [49] construct and simulate anatomical skeletal-muscle geometry. Active Volumetric Musculoskeletal Systems [6] couples deformable muscles in close contact to multibody dynamics. Sueda et al. [47] and Sachdeva et al. [41] address tendon routing and transmission in detailed hand mod-

els. Lee et al. [20] integrate volumetric actuation into dexterous motion control and trajectory optimization. These works provide direct precedents for tissue-to-bone force feedback. Projective dynamics [1] and XPBD [31] provide established deformation solvers. Romeo et al. [40] develop anisotropic muscle-fiber and superficial-fascia constraints within XPBD for visual effects animation; MGPBD [24] uses algebraic multigrid to accelerate a global XPBD solve. RLMuscle [26] instead uses geometric multigrid for volumetric muscles and learns activations on a separate one-dimensional surrogate. Its measured EMG prior helps select among redundant activation solutions, while the three-dimensional stage produces the resulting deformation. Our formulation studies model-based RL with quasi-static deformable bellies, simplified active-fiber energy, and endpoint/wrapping reactions inside skeletal dynamics. Bone pose changes the attachment and contact energies; their reactions supply skeletal loads. Virtual-work and energy-gradient tests assess numerical consistency alongside learned tracking.

3. Volumetric Muscle Control

We learn motion tracking with volumetric muscles using the pipeline in Figure 1. A policy controls muscle activation to generate forces that drive skeletal motion. Skeletal pose q changes attachment and wrapping energies, altering muscle equilibrium and the resulting forces. We train a differentiable world model on these transitions and use its predictions of future motion to improve the tracking policy.

3.1. Motion-tracking objective

The reference motion is the tracking target, represented by g_t at control step t. The policy $u_t = \pi_\theta(O(s_t), g_t)$ maps the physical-state observation $O(s_t)$ and reference input g_t to muscle excitation vector u_t. Here s_t is the simulator state vector, O is the observation function, and θ denotes policy parameters. We optimize

$$\max_\theta \mathbb{E}_{\pi_\theta, \mathcal{F}_{\text{vol}}} \left[\sum_{t=0}^{T-1} \gamma^t r(s_t, g_t) \right], \quad s_{t+1} = \mathcal{F}_{\text{vol}}(s_t, u_t). \tag{1}$$

Here t indexes control steps, T is the rollout length, γ is the discount factor, and r measures agreement with the reference. $\mathcal{F}_{\text{vol}}$ is the volumetric muscle simulation engine. State s_t includes skeletal positions q_t, velocities v_t, and muscle activations a_t.

3.2. Model-based learning of motion control

We follow the model-based RL paradigm [7, 8, 27, 58], and extend it to motion control with volumetric muscles. We run the policy in the volumetric muscle simulator and store

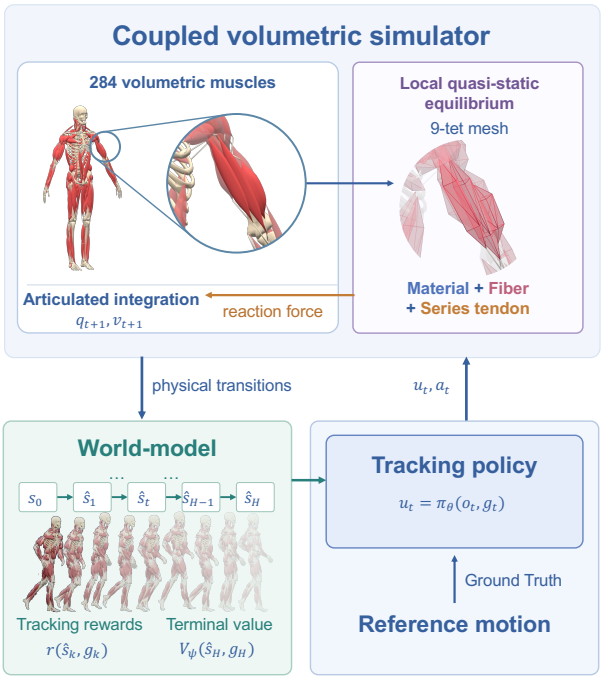

Figure 1. **Volumetric muscle control.** Skeletal pose changes muscle attachment and wrapping geometry; quasi-static tissue equilibrium generates forces that drive the skeleton. Physical simulator transitions train the world model; its predictions are used to improve the tracking policy.

the resulting state transitions in a replay buffer. The world model predicts the next state from the current state and muscle excitation:

$$\hat{s}_{j+1} = f_\omega(\hat{s}_j, u_j), \tag{2}$$

where f_ω is the transition model with parameters ω, $\hat{s}_j$ is the predicted state at rollout step j, and u_j is the muscle excitation vector. Initial state $\hat{s}_0$ comes from physical replay. Model fitting unrolls recorded commands and minimizes

$$\mathcal{L}_{\text{WM}} = \mathbb{E}_{\mathcal{B}} \sum_{j=1}^{K} \|\hat{s}_j - s_j\|_W^2 . \tag{3}$$

The replay buffer $\mathcal{B}$ supplies sequences of K transitions; s_j is the recorded state vector, and positive-semidefinite matrix W weights the squared state-prediction errors. To improve the policy, we predict future states with the world model using commands from the current policy and maximize

$$J(\theta) = \mathbb{E}\Big[\sum_{j=0}^{H-1} \gamma^j r(\hat{s}_j, g_j) + \gamma^H V_\psi(\hat{s}_H, g_H) \Big], \tag{4}$$

where H is the prediction horizon, θ is the policy parameter vector, and critic V_ψ, with parameters ψ, predicts terminal tracking return. Reward r, reference input g_j, and discount γ retain their tracking meanings. The expectation averages

trajectories predicted from initial states sampled from the replay buffer.

Training alternates between fitting the world model to simulator transitions and updating the policy through the fitted world model.

3.3. Two-way volumetric muscle dynamics

Muscle forces drive skeletal motion, while skeletal motion changes muscle energy, deformation, and force transmission. We update muscle attachments and wrapping geometry from skeletal pose q. Policy commands change muscle activation and thus the active contraction energy within the tissue. For the current pose and activations, we solve for the equilibrium tissue shape and apply the resulting attachment and contact forces to the bones, completing the two-way coupling.

Muscle equilibrium. We represent each muscle in our human model as a deformable tetrahedral mesh. At each muscle update time t, we hold skeletal pose and activation fixed and find the vertex positions that minimize the muscle energy. For brevity, we omit the time subscript t and the muscle index on individual energy components below:

$$\begin{aligned} x_m^* &\approx \arg\min_{x_m} E_m(x_m; q, a_m), \\ E_m &= E_{\text{mat}} + E_{\text{act}} + E_{\text{att}} \\ &\quad + E_{\text{con}} + E_{\text{bar}}, \qquad E = \sum_m E_m, \end{aligned} \tag{5}$$

where m indexes muscles, x_m collects the world-space positions of the vertices in muscle m's mesh, x_m^* gives their approximate equilibrium positions, and a_m is muscle activation. E_{mat} is the passive material energy, E_{act} the active contraction energy, E_{att} the endpoint attachment energy, E_{con} the wrapping contact energy, and E_{bar} a volume barrier against element collapse. The total energy E is the sum of the per-muscle energies E_m.

Passive and active energies. For each tetrahedron, we compute the deformation gradient

$$F_e = D_{s,e} D_{0,e}^{-1}, \qquad J_e = \det F_e, \tag{6}$$

where e indexes tetrahedra, and $D_{s,e}$ and $D_{0,e}$ contain three edge vectors in the deformed and rest configurations. Thus F_e describes local deformation, while J_e is the deformed-to-rest volume ratio. The passive energy is

$$E_{\text{mat}} = \sum_e \frac{w_e}{2} \left(\|F_e - R_e\|_F^2 + \nu \|F_e - Z_e\|_F^2 \right), \tag{7}$$

where w_e is a material weight incorporating the tetrahedron's rest volume, ν controls the volume penalty relative to the shape penalty, and $\|\cdot\|_F$ is the Frobenius norm. The rotation R_e and volume-preserving projection Z_e are computed from F_e using projective dynamics (PD) [1, 20]. The two energy terms penalize shape distortion and volume change, respectively.

We use a quadratic fiber energy for active contraction [20]:

$$E_{\text{act}} = \sum_e \frac{a_m c_e}{2} \|F_e d_e\|^2. \tag{8}$$

Here a_m is the activation of muscle m, d_e is the unit rest-fiber direction, and c_e is active stiffness times rest volume. The length $\|F_e d_e\|$ is fiber stretch. Increasing activation increases the tendency to shorten fibers; the achieved shape balances this contraction against material, attachment, and contact forces. The volume barrier E_{bar} discourages flattened or inverted tetrahedra during the equilibrium solve.

Muscle attachments. Skeletal pose q enters the muscle energy through the endpoint connections and wrapping geometry illustrated in Figure 3. For endpoint h, routed-tendon energy $E_{\text{route},h}$ models a tendon following bone-fixed waypoints; direct-attachment energy $E_{\text{direct},h}$ models a spring to a bone:

$$\begin{aligned} L_h &= \|x_{i_h} - y_{h,1}(q)\| \\ &\quad + \sum_{j=1}^{n_h - 1} \|y_{h,j+1}(q) - y_{h,j}(q)\|, \\ E_{\text{route},h} &= \tfrac{1}{2} k_h [L_h - L_{0,h}]_+^2, \\ E_{\text{direct},h} &= \tfrac{1}{2} k_h \|x_{i_h} - y_h(q)\|^2. \end{aligned} \tag{9}$$

Here i_h identifies the muscle-tip vertex x_{i_h} shown in Figure 2(b), and $y_{h,j}(q)$ are n_h ordered bone-fixed waypoints. Length L_h includes the tip-to-first-waypoint segment and the remaining route; $L_{0,h}$ is slack length, k_h is connection stiffness, and $[z]_+ = \max(z, 0)$ is the positive part. Point $y_h(q)$ is the direct bone attachment. The routed tendon pulls only when taut; the direct spring is bilateral. Each end selects one connection type, giving

$$E_{\text{att}} = \sum_{h \in \mathcal{R}_m} E_{\text{route},h} + \sum_{h \in \mathcal{D}_m} E_{\text{direct},h}. \tag{10}$$

Here $\mathcal{R}_m$ and $\mathcal{D}_m$ contain the routed and directly attached endpoints of muscle m, respectively. These disjoint sets partition its two endpoints, so each contributes through its configured connection.

Wrapping contact. In addition to endpoint connections, muscle bellies interact with bone-carried cylinders that guide wrapping around joints. The contact energy is

$$E_{\text{con}} = \sum_{s \in \mathcal{S}} \tfrac{1}{2} k_c [R_s - \rho_s(x, q)]_+^2, \tag{11}$$

where x contains tissue vertices, q is skeletal pose, and $\mathcal{S}$ is the set of checked belly-edge and endpoint-connection segments. Parameter k_c is contact stiffness, R_s is the cylinder radius including contact clearance, and ρ_s is the minimum radial distance between segment s and its cylinder axis. Both bone-side segment endpoints and cylinder poses depend on q. The contact penalty produces equal and opposite reactions on tissue and bones.

Muscle solver. We solve the equilibrium problem in Eq. 5 using PD local/global updates of the vertex positions, with updated attachment/contact terms and line search. We monitor convergence using nodal force residuals:

$$f_{m,i} = -\frac{\partial E_m}{\partial x_{m,i}}, \qquad r_m = \max_i \|f_{m,i}\|_2 \le \varepsilon. \tag{12}$$

Here i indexes muscle vertices, $f_{m,i}$ is nodal force, r_m is the largest nodal residual, and $\varepsilon = 10^{-3} F_0$ is the stopping threshold. The force scale is $F_0 = \max_m F_{\max,m}$, where $F_{\max,m}$ denotes the maximum force of muscle m.

Skeletal forces. At the computed equilibrium, we assemble the attachment and contact reactions into a generalized skeletal force:

$$\begin{aligned} Q^{\text{mus}} &= -\left.\frac{\partial E(x;q,a)}{\partial q}\right|_{x=x^*} \\ &= \sum_b \left(J_{v,b}^\top F_b + J_{\omega,b}^\top \tau_b\right), \end{aligned} \tag{13}$$

where x^* collects the muscle equilibria, a collects activations, and Q^{mus} is generalized muscle force. The derivative holds tissue positions and activation fixed and denotes skeletal virtual work. For bone b, F_b and τ_b are resultant attachment/contact force and torque about its center of mass. Jacobians $J_{v,b}$ and $J_{\omega,b}$ map skeletal velocity to that body's translational and angular velocities.

Skeletal integration. We advance the skeleton with the computed muscle load, holding it fixed between muscle updates:

$$(q_{n+1}, v_{n+1}) = \mathcal{D}_{\Delta t}(q_n, v_n; Q_k^{\text{mus}}). \tag{14}$$

Here n indexes rigid steps, k indexes the enclosing muscle-update interval, and Q_k^{mus} is held throughout that interval. The map $\mathcal{D}_{\Delta t}$ advances skeletal positions q_n and velocities v_n by rigid step size Δt, including inertia, gravity, and external contacts. The next muscle solve uses the updated bone pose, closing the two-way loop.

4. Experiments

We evaluate VolMusCo on a large motion set, measuring motion-tracking accuracy, mechanical consistency, and computational efficiency. The experiments separate actuation-model choice, deformation-dependent force feedback, mesh resolution, and world-model learning efficiency.

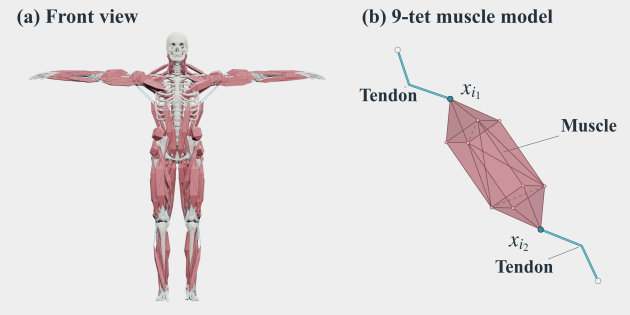

Figure 2. **Coarse full-body muscle geometry.** (a) Full-body geometry comprising 284 muscles and 2,556 tetrahedra. (b) A 9-tet muscle mesh with 10 vertices and tendon connections at the tips $\mathbf{x}_{i_1}$ and $\mathbf{x}_{i_2}$.

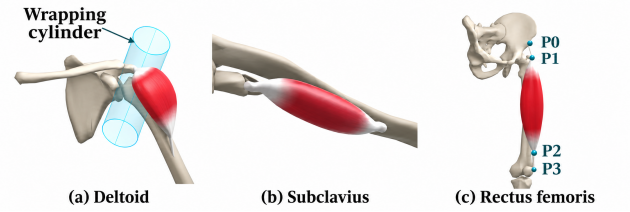

Figure 3. **Muscle–joint interactions.** We use 3 modeling configurations according to how muscle bellies and tendons interact with joints. (a) The deltoid belly wraps around the shoulder joint; we model its wrapping contact with a cylinder. (b) For the subclavius, belly wrapping has little effect on joint motion, so we use direct bone attachment. (c) The rectus femoris has long, slender tendons whose paths determine how muscle forces act on the joints; we represent these routes with bone-fixed waypoints labeled P_i.

4.1. Assets and training configurations

We construct the high-resolution muscle meshes in Figure 3 by sweeping cross-sections along MASS Hill-type muscle paths [21] and fitting their profiles and target volumes to RLMuscle anatomical meshes [26].

For RL, we fit the 10-vertex, 9-tet mesh shown in Figure 2(b) to each high-resolution muscle belly. We select 2 points on the Hill path to define the belly endpoints and constrain the 2 tip vertices to these positions during mesh construction. The remaining 8 vertices form 2 quadrilateral cross-sections, whose positions are optimized to match the target volume and approximate the reference belly shape while retaining the 9-tet connectivity. This fitting accommodates the curved or flattened profiles of selected shoulder and gluteal bellies. Axial reference fibers and the tendon paths outside the belly are retained. The model contains 284 muscle units, 2,840 vertices, and 2,556 tetrahedra.

The motion references combine 60 LaFAN1 [15] clips from 5 subjects, custom keyframed lifting exercises and

object-manipulation sequences, and motions reconstructed from monocular videos. We use the body-landmark estimator from BlazePose GHUM Holistic [12], implemented as MediaPipe Pose Full, to estimate 2D image landmarks and pelvis-relative 3D landmarks. We adapt these landmarks to our skeleton's bone lengths, recover joint angles with joint-constrained inverse kinematics, and apply temporal smoothing.

We use MuJoCo (MJWarp) as the simulation engine and solve tissue equilibrium in Warp, running 2,048 parallel environments on a single NVIDIA RTX 4090 GPU (24 GB). We calibrate material and connection parameters per muscle. The 3 muscle–joint configurations in Figure 3 are wrapping contact with bone-carried cylinders (a), direct bone attachment (b), and tendon routing through bone-fixed waypoints (c).

4.2. Local solve and multi-rate collection

We run the controller at 50 Hz, solve muscle equilibrium at 100 Hz, and integrate skeletal dynamics at 1 kHz. Each 20 ms control interval contains 2 muscle solves and 20 rigid steps. We hold the generalized muscle load for 10 rigid steps after each solve and use the updated skeletal pose for the next solve.

PD uses local/global updates with line search and a maximum nodal force residual threshold $\varepsilon = 10^{-3}F_0$, where $F_0 = \max_m F_{\max,m}$ and $F_{\max,m}$ is the maximum force of muscle m. The volume barrier enforces a lower threshold $J_{\min} = 0.1$ on the local deformed-to-rest volume ratio $J_e = \det F_e$. Muscle and wrapping solves are capped at 600 and 50 iterations, respectively. For each of 240 independently initialized poses sampled from 60 LaFAN1 clips, we solve all 284 muscles at 3 activation levels: 0.005, 0.5, and 1.0. All $3 \times 240 \times 284 = 204{,}480$ muscle solves satisfy the residual threshold. We reject a candidate pose if any muscle solve fails to converge.

4.3. Main results

On LaFAN1 clips, VolMusCo achieves a tracking success rate of 97.9% and a mean local-body position error of 5.52 cm within 10 h of training. The simulator collects 815 physical control transitions per second in aggregate, with a total training cost of 16 GPU-hours. Figure 4 presents tracking examples from LaFAN1, custom fitness motions, and motions reconstructed from monocular videos. Each row follows a single closed-loop rollout at aligned timestamps. Cameras and activation color scales are fixed within each sequence. Video demonstrations and interactive visualizations are available at https://inckieultra.github.io/VolMusCo-Project-Page/.

Table 1. **Mesh-resolution ablation.** We compare 9-tet and 230-tet muscle meshes. Δ_T is the maximum relative L_2 discrepancy in tendon tensions over 4 static conditions, normalized by the 9-tet result; direct-attachment and wrapping-contact forces are excluded. Δ_q is the pooled joint-angle RMSE over 50 control timestamps and 50 scalar joints in a 1 s full-body rollout, excluding the free root. Samples/s reports physical control transitions per second for the coarse and fine meshes, respectively.

Mesh pair	Tets/unit	Samples/s	Δ_T (%)	Δ_q (°)
Coarse / fine	9 / 230	815 / 0.05	9.46	1.87

4.4. Mesh-resolution ablation

We compare the high-resolution 230-tet muscle meshes illustrated in Figure 3 with the 9-tet meshes in Figure 2(b) to evaluate simulation fidelity and RL collection cost. Both resolutions use the same material/active energies, endpoint connections, wrapping geometry, rest volume, axial length, fiber directions, physical parameters, and stopping targets.

Table 1 reports relative tension discrepancy, scripted joint-motion discrepancy, and aggregate physical control transitions per second at a fixed environment count. We use matched activation/path sweeps and short free-joint rollouts to measure simulation fidelity. Timing includes tissue solves, attachment/contact reaction assembly, and rigid stepping.

4.5. Force transmission and mechanical checks

We sweep activation and skeletal pose, recording fiber stretch, routed-tendon tension, direct-attachment forces, contact penetration, and generalized torque. Fixed-shape finite differences of attachment/contact energy check the reaction map. Slack-load tests disable direct attachments and contact, isolating the routed tendon; mixed-connection tests include every reaction channel. The maximum isolated slack-tendon load is 0 N and generalized-load error is 0.0441% (rotational) and 1.31% (translational).

We assess the complete reaction-to-generalized-force map through virtual power. A separate reduced-energy test resolves local equilibrium after each pose perturbation. The relative virtual-power discrepancy is $8.87 \times 10^{-7}\%$ and the free-motion energy drift is 23.1%.

4.6. Deformation-feedback ablation

To isolate deformation-dependent force feedback, we compare VolMusCo with a frozen nominal map $\widetilde{Q}^{\mathrm{mus}}(q, a)$. We fit a separate MLP with 2 hidden layers of width 64 for each muscle to equilibrium reactions from separate pose/activation samples at nominal material parameters. The networks predict attachment, tendon, and wrapping reactions from current geometry and activation, which are assembled into skeletal loads. Network weights and nominal material response remain fixed, while inputs and

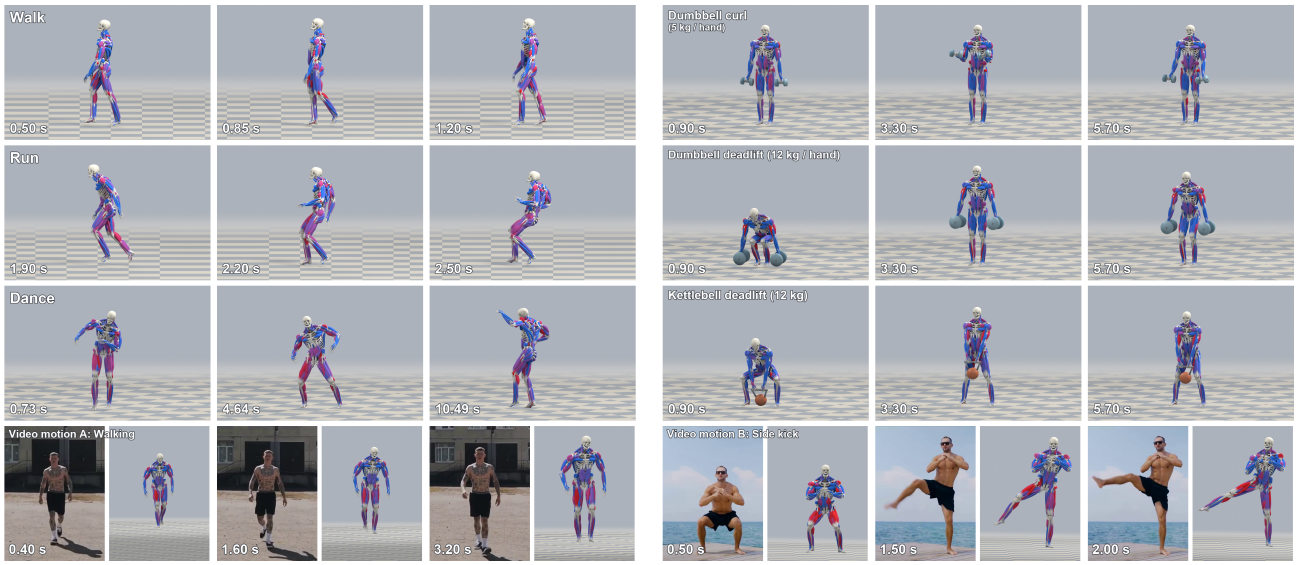

Figure 4. **Volumetric motion tracking.** The rows show a LaFAN1 motion, a custom fitness motion, and a motion reconstructed from a monocular video. Muscle color indicates activation level, ranging from blue (zero activation) to red (full activation).

Table 2. **Tissue-response feedback under material change.** A fixed policy is evaluated after doubling Young's modulus. Paired values follow the Ours / Frozen order. Δ_Q measures frozen-load discrepancy relative to Ours.

Change	Method	Δ_Q (%)	Local error (cm)	Success (%)
$E \times 2$	Ours / Frozen	11.48	9.71 / 8.89	68.0 / 76.1

loads update with skeletal motion; no tissue equilibrium solve is performed for the map.

In Figure 5, we simulate the same recorded aiming activation sequence from identical initial skeletal states. Panel (a) compares VolMusCo at nominal stiffness, coupled simulation with doubled Young's modulus ($E \times 2$), and the frozen nominal map. Active strength and tendon parameters are unchanged. The frozen-map rollout stops at 0.91 s when its inputs exceed the calibrated range. Panel (b) compares coarse 9-tet and fine 230-tet meshes under the same activations. Table 2 reports fixed-policy tracking after controlled belly-shape or material changes, with strength scales, tendon parameters, skeleton, routing, contacts, observations, and goals held fixed.

4.7. Independent actuation-model ablation

Table 3 compares Hill-type and coupled volumetric actuation. Each controller and world model is trained independently using the same model-based learning procedure and network architectures. We match the skeleton, actuator correspondence, nominal strength, motion references, rewards, initial-state distribution, control rate, policy observations, and physical-transition and optimization-update budgets. Wall-clock cost is reported separately. Backend-specific internal states, including volumetric nodal coordinates, are excluded from policy and world-model inputs.

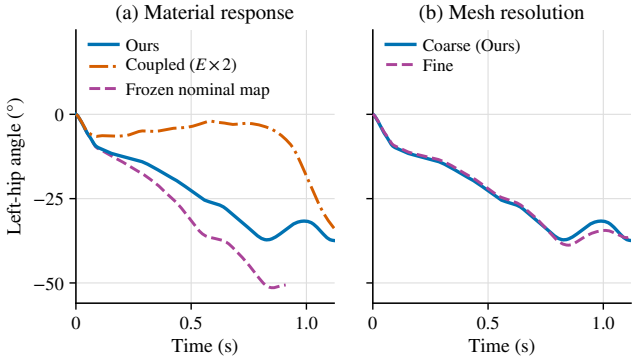

Figure 5. **Simulation with identical activation sequences.** Left-hip angles from matched initial skeletal states: (a) VolMusCo at nominal stiffness, coupled simulation with doubled Young's modulus, and the frozen nominal map; (b) coarse 9-tet versus fine 230-tet meshes. The frozen-map rollout stops at 0.91 s when its inputs exceed the calibrated range.

Table 3. **Actuation-model ablation.** Both backends are trained independently with matched learning/evaluation conditions. Local error is in cm, success in percent, and throughput in aggregate physical control transitions/s.

Actuation	Local error	Success	Samples/s
Hill line	4.92	97.9%	9030
Volumetric	5.52	97.9%	815

We evaluate tracking on 47 clips from the LAFAN1 motion dataset, initializing all compared methods from the same reference start states and using deterministic policy rollouts. Metrics include root, local-body, and endpoint errors, tracking success rate, normalized return, activation

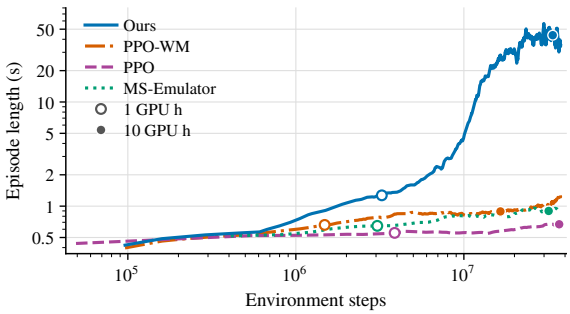

Figure 6. **Learning efficiency in the volumetric simulator.** Mean episode length is reported in seconds. Transitions (millions) and training hours to reach a mean episode length of $\tau = 10\,\mathrm{s}$ are reported separately. Open and filled dots mark where each method had used 1 and 10 GPU hours.

magnitude and variation, and training cost. An episode terminates early if the absolute root-height error exceeds $0.25\,\mathrm{m}$, the root-tilt discrepancy $|R_{zz} - R_{zz}^{\mathrm{ref}}|$ exceeds 0.8, or horizontal root-position error exceeds $10\,\mathrm{m}$. A rollout is considered successful if it reaches $10\,\mathrm{s}$ or the end of the reference clip, whichever occurs first, without triggering a failure termination.

4.8. World-model learning efficiency

Table 4 compares our world-model training at critic horizons of 16, 32, and 48 steps with three baselines: PPO [44]; MS-Emulator [54], model-free adversarial imitation with a value-guided flow policy; and PPO-WM, which keeps our world model, imagined rollouts, rewards, and critic but updates the policy with PPO. All variants use the same motion split, reward, common observations, policy/critic family, and physical-transition budget.

At the default 32-step horizon, mean training episode length reaches 38.2 s with a normalized per-step reward of 0.709. Figure 6 plots mean episode length against physical transitions and training time, including collection and all learning updates. We report transitions and hours to reach a mean episode length of 10 s; runs below this threshold are censored at their budget. On a shared fixed replay set with identical evaluation normalization, four-step world-model state prediction MSE is 5.283, 3.798, and 3.688 for horizons of 16, 32, and 48 steps, respectively.

4.9. Muscle recruitment

We compare activations during LaFAN1 walking with surface EMG [34]. Each of the 12 walking clips is played from its first frame for up to 180 s; when the character falls, it restarts from the reference state at that frame, and the first second after each restart is discarded, which leaves 97 right-leg strides. For each electrode we take the force-weighted mean activation of the synergist group under it (Figure 7). Without EMG supervision, the coarse timing of recruitment emerges: the quadriceps is most active in early

Table 4. **Learning efficiency in the volumetric simulator.** Mean episode length is reported in control steps. Transitions (millions) and training hours to reach a mean episode length of $\tau = 10\,\mathrm{s}$ are reported separately. World-model labels indicate critic horizons in control steps.

Variant	Ep. length	Trans.@τ	Hours@τ
PPO	31.6	>188.7	>50.2
Critic horizon: 16	503.4	11.616	3.64
Critic horizon: 32	506.5	9.632	3.00
Critic horizon: 48	506.5	11.487	3.56

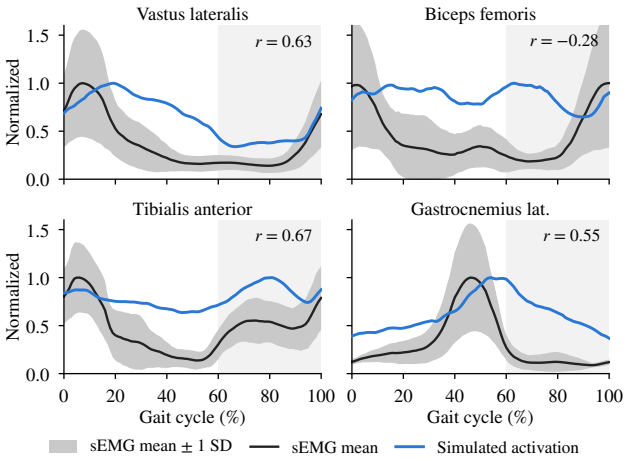

Figure 7. Simulated activation (blue, mean of 97 strides) versus sEMG (black: mean; gray: ±1 SD; 13 adults at 4 km/h [34]) over the gait cycle, right leg, peak-normalized; 0% is heel strike, shading marks typical swing (60–100%). r: Pearson correlation of the mean curves.

stance ($r = 0.63$), the triceps surae in push-off ($r = 0.55$, 11% of the cycle late), and the tibialis anterior follows the stance-to-swing pattern ($r = 0.67$), though without its burst after heel strike. The modulation does not: between bursts the simulated groups stay at 34–65% of their peak, against 9–19% in sEMG, so only 26% of the cycle lies within one standard deviation of the human mean, and the hamstrings remain active through stance. For the single muscle under each electrode, the mean correlation is $r = 0.48$. Kinematic tracking alone does not determine how work is split among synergists.

4.10. Limitations

The coarse belly approximates anatomy. Direct endpoints are spring approximations, while segment/cylinder wrapping permits finite penetration and omits explicit fascia and general tissue–tissue contact. Quasi-static tissue omits belly inertia, oscillations, and impulse transmission; the fiber response omits force–velocity effects. Finite equilibrium residuals and held-force stepping can affect work accuracy. Motion tracking alone does not resolve muscle-recruitment

ambiguity or establish physiological validity [17].

5. Conclusion

We present VolMusCo, a model-based reinforcement learning framework for full-body motion control with volumetric muscles and two-way muscle–skeleton coupling. Bone pose changes attachment and wrapping energy, while tissue equilibrium supplies skeletal reactions. Simplified tetrahedral muscle meshes make volumetric simulation efficient enough for learning. Our method achieves accurate full-body motion tracking on LaFAN1, custom fitness sequences, and motions reconstructed from videos.

References

[1] Sofien Bouaziz, Sebastian Martin, Tiantian Liu, Ladislav Kavan, and Mark Pauly. Projective dynamics: Fusing constraint projections for fast simulation. *ACM Transactions on Graphics*, 33(4):154:1–154:11, 2014. 3, 4

[2] Vittorio Caggiano, Huawei Wang, Guillaume Durandau, Massimo Sartori, and Vikash Kumar. MyoSuite: A contact-rich simulation suite for musculoskeletal motor control. In *Learning for Dynamics and Control Conference*, pages 492–507, 2022. 2

[3] Mia Chiquier and Carl Vondrick. Muscles in action. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22091–22101, 2023. 2

[4] Scott L. Delp, Frank C. Anderson, Allison S. Arnold, Peter Loan, Ayman Habib, Chand T. John, Eran Guendelman, and Darryl G. Thelen. OpenSim: Open-source software to create and analyze dynamic simulations of movement. *IEEE Transactions on Biomedical Engineering*, 54 (11):1940–1950, 2007. 2

[5] Christopher L. Dembia, Nicholas A. Bianco, Antoine Falisse, Jennifer L. Hicks, and Scott L. Delp. OpenSim Moco: Musculoskeletal optimal control. *PLOS Computational Biology*, 16(12):e1008493, 2020. 2

[6] Ye Fan, Joshua Litven, and Dinesh K. Pai. Active volumetric musculoskeletal systems. *ACM Transactions on Graphics*, 33(4), 2014. 1, 2

[7] Yusen Feng, Xiyan Xu, and Libin Liu. MuscleVAE: Model-based controllers of muscle-actuated characters. In *SIGGRAPH Asia Conference Papers*, 2023. 2, 3

[8] Levi Fussell, Kevin Bergamin, and Daniel Holden. SuperTrack: Motion tracking for physically simulated characters using supervised learning. *ACM Transactions on Graphics*, 40(6), 2021. 2, 3

[9] Thomas Geijtenbeek, Michiel van de Panne, and A. Frank van der Stappen. Flexible muscle-based locomotion for bipedal creatures. *ACM Transactions on Graphics*, 32(6), 2013. 2

[10] Ignat Georgiev, Krishnan Srinivasan, Jie Xu, Eric Heiden, and Animesh Garg. Adaptive horizon actor-critic for policy learning in contact-rich differentiable simulation. In *ICML*, pages 15418–15437, 2024. 2

[11] Hartmut Geyer and Hugh Herr. A muscle-reflex model that encodes principles of legged mechanics produces human walking dynamics and muscle activities. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 18(3): 263–273, 2010. 2

[12] Ivan Grishchenko, Valentin Bazarevsky, Andrei Zanfir, Eduard Gabriel Bazavan, Mihai Zanfir, Richard Yee, Karthik Raveendran, Matsvei Zhdanovich, Matthias Grundmann, and Cristian Sminchisescu. BlazePose GHUM Holistic: Real-time 3D human landmarks and pose estimation. arXiv:2206.11678, 2022. 6

[13] Michael Günther, Oliver Röhrle, Daniel F. B. Haeufle, and Syn Schmitt. Spreading out muscle mass within a Hill-type model: A computer simulation study. *Computational and Mathematical Methods in Medicine*, 2012:848630, 2012. 1

[14] Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse control tasks through world models. *Nature*, 640:647–653, 2025. 2

[15] Félix G. Harvey, Mike Yurick, Derek Nowrouzezahrai, and Christopher Pal. Robust motion in-betweening. *ACM Transactions on Graphics*, 39(4):60:1–60:12, 2020. 2, 5

[16] Kaibo He, Chenhui Zuo, Chengtian Ma, and Yanan Sui. DynSyn: Dynamical synergistic representation for efficient learning and control in overactuated embodied systems. In *ICML*, pages 18115–18132, 2024. 2

[17] 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):020905, 2015. 9

[18] Archibald 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

[19] Michael Janner, Justin Fu, Marvin Zhang, and Sergey Levine. When to trust your model: Model-based policy optimization. In *Advances in Neural Information Processing Systems*, 2019. 2

[20] Seunghwan Lee, Ri Yu, Jungnam Park, Mridul Aanjaneya, Eftychios Sifakis, and Jehee Lee. Dexterous manipulation and control with volumetric muscles. *ACM Transactions on Graphics*, 37(4):57:1–57:13, 2018. 1, 3, 4

[21] Seunghwan Lee, Moonseok Park, Kyoungmin Lee, and Jehee Lee. Scalable muscle-actuated human simulation and control. *ACM Transactions on Graphics*, 38(4):73:1–73:13, 2019. 1, 2, 5

[22] Sung-Hee Lee, Eftychios Sifakis, and Demetri Terzopoulos. Comprehensive biomechanical modeling and simulation of the upper body. *ACM Trans. Graph.*, 28(4), 2009. 1

[23] Yoonsang Lee, Moon Seok Park, Taesoo Kwon, and Jehee Lee. Locomotion control for many-muscle humanoids. *ACM Transactions on Graphics*, 33(6), 2014. 2

[24] Chunlei Li, Peng Yu, Tiantian Liu, Siyuan Yu, Yuting Xiao, Shuai Li, Aimin Hao, Yang Gao, and Qinping Zhao. MGPBD: A multigrid accelerated global XPBD solver. In *ACM SIGGRAPH Conference Papers*, 2025. 3

[25] Chengkun Li, Cheryl Wang, Bianca Ziliotto, Merkourios Simos, Jozsef Kovecses, Guillaume Durandau, and Alexander Mathis. Towards embodied AI with MuscleMimic: Unlocking full-body musculoskeletal motor learning at scale. arXiv:2603.25544, 2026. 1, 2

[26] Chunlei Li, Siyuan Yu, Yang Gao, Shuai Li, Peng Yu, Aimin Hao, and He Wang. RLMuscle: Muscle-driven character animation via reinforcement learning. arXiv:2610.30033, 2026. 1, 3, 5

[27] Qiayuan Liao, Yiyang Shao, Junfeng Long, Kevin Zakka, Roberto Horowitz, Pieter Abbeel, and Koushil Sreenath. SuperMimic: Scaling humanoid whole-body control with one differentiable world model. In *Conference on Robot Learning*, 2026. 2, 3

[28] Qiayuan Liao, Takara E. Truong, Xiaoyu Huang, Yuman Gao, Guy Tevet, Koushil Sreenath, and C. Karen Liu. BeyondMimic: From motion tracking to versatile humanoid control via guided diffusion. *Science Robotics*, 11(117): eadx8924, 2026. 2

[29] Zhengyi Luo, Jinkun Cao, Josh Merel, Alexander Winkler, Jing Huang, Kris Kitani, and Weipeng Xu. Universal humanoid motion representations for physics-based control. In *ICLR*, 2024. 2

[30] Zhengyi Luo, Ye Yuan, Tingwu Wang, Chenran Li, Fernando Castañeda, Sirui Chen, Zi-Ang Cao, Jiefeng Li, David Minor, Qingwei Ben, Jinhyung Park, David Sami, Zi Wang, Xingye Da, Runyu Ding, Cyrus Hogg, Lina Song, Edy Lim, Eugene Jeong, Tairan He, Haoru Xue, Wenli Xiao, Simon Yuen, Jan Kautz, Yan Chang, Umar Iqbal, Linxi Fan, and Yuke Zhu. SONIC: Supersizing motion tracking for natural humanoid whole-body control. *Science Robotics*, 11(117): eaed4592, 2026. 2

[31] Miles Macklin, Matthias Müller, and Nuttapong Chentanez. XPBD: Position-based simulation of compliant constrained dynamics. In *International Conference on Motion in Games*, pages 49–54, 2016. 3

[32] Aleka McAdams, Yongning Zhu, Andrew Selle, Mark Empey, Rasmus Tamstorf, Joseph Teran, and Eftychios Sifakis. Efficient elasticity for character skinning with contact and collisions. *ACM Trans. Graph.*, 30(4), 2011. 1

[33] Matthew Millard, Thomas Uchida, Ajay Seth, and Scott L. Delp. Flexing computational muscle: Modeling and simulation of musculotendon dynamics. *Journal of Biomechanical Engineering*, 135(2):021005, 2013. 1

[34] Luís Moreira, Joana Figueiredo, Pedro Fonseca, João P. Vilas-Boas, and Cristina P. Santos. Lower limb kinematic, kinetic, and EMG data from young healthy humans during walking at controlled speeds. *Scientific Data*, 8:103, 2021. 8

[35] Jungnam Park, Sehee Min, Phil Sik Chang, Jaedong Lee, Moon Seok Park, and Jehee Lee. Generative GaitNet. In *ACM SIGGRAPH Conference Papers*, 2022. 2

[36] Jungnam Park, Moon Seok Park, Jehee Lee, and Jungdam Won. Bidirectional GaitNet: A bidirectional prediction model of human gait and anatomical conditions. In *ACM SIGGRAPH Conference Papers*, 2023. 2

[37] Jungnam Park, Euikyun Jung, Jehee Lee, and Jungdam Won. MAGNET: Muscle activation generation networks for diverse human movement. In *ACM SIGGRAPH Conference Papers*, 2025. 1, 2

[38] 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*, 37(4), 2018. 2

[39] 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. 2

[40] Marco Romeo, Carlos Monteagudo, and Daniel Sánchez-Quirós. Muscle and fascia simulation with extended position based dynamics. *Computer Graphics Forum*, 39(1): 134–146, 2020. 3

[41] Prashant Sachdeva, Shinjiro Sueda, Susanne Bradley, Mikhail Fain, and Dinesh K. Pai. Biomechanical simulation and control of hands and tendinous systems. *ACM Transactions on Graphics*, 34(4), 2015. 2

[42] David Schneider, Simon Reiß, Marco Kugler, Alexander Jaus, Kunyu Peng, Susanne Sutschet, M. Saquib Sarfraz, Sven Matthiesen, and Rainer Stiefelhagen. Muscles in time: Learning to understand human motion in-depth by simulating muscle activations. In *Advances in Neural Information Processing Systems*, 2024. 2

[43] John Schulman, Philipp Moritz, Sergey Levine, Michael I. Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. In *ICLR*, 2016.

[44] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. arXiv:1707.06347, 2017. 8

[45] Pierre Schumacher, Daniel Häufle, Dieter Büchler, Syn Schmitt, and Georg Martius. DEP-RL: Embodied exploration for reinforcement learning in overactuated and musculoskeletal systems. In *ICLR*, 2023. 2

[46] Merkourios Simos, Alberto Chiappa, and Alexander Mathis. KINESIS: Motion imitation for human musculoskeletal locomotion. In *2026 IEEE International Conference on Robotics and Automation (ICRA)*, pages 20907–20914. IEEE, 2026. 2

[47] Shinjiro Sueda, Andrew Kaufman, and Dinesh K. Pai. Musculotendon simulation for hand animation. *ACM Transactions on Graphics*, 27(3), 2008. 2

[48] Richard S. Sutton. Dyna, an integrated architecture for learning, planning, and reacting. *ACM SIGART Bulletin*, 2(4): 160–163, 1991. 2

[49] Joseph Teran, Eftychios Sifakis, Silvia S. Blemker, Victor Ng-Thow-Hing, Cynthia Lau, and Ronald Fedkiw. Creating and simulating skeletal muscle from the visible human data set. *IEEE Transactions on Visualization and Computer Graphics*, 11(3):317–328, 2005. 1, 2

[50] Chen Tessler, Yunrong Guo, Ofir Nabati, Gal Chechik, and Xue Bin Peng. MaskedMimic: Unified physics-based character control through masked motion inpainting. *ACM Transactions on Graphics*, 43(6), 2024. 2

[51] Jack M. Wang, Samuel R. Hamner, Scott L. Delp, and Vladlen Koltun. Optimizing locomotion controllers using

biologically-based actuators and objectives. *ACM Transactions on Graphics*, 31(4), 2012. 2

[52] William Wang, Nicholas Bianco, Guy Tevet, Jennifer Hicks, C. Karen Liu, Scott Delp, and Kayvon Fatahalian. Learning realistic athletic sprinting without demonstrations. arXiv:2609.11083, 2026. 2

[53] Ying Wang, Jasper Verheul, Sang-Hoon Yeo, Nima Khademi Kalantari, and Shinjiro Sueda. Differentiable simulation of inertial musculotendons. *ACM Transactions on Graphics*, 41 (6):272:1–272:11, 2022. 1

[54] Yunyue Wei, Chenhui Zuo, Shanning Zhuang, Haixin Gong, Yaming Liu, and Yanan Sui. Scaling whole-body human musculoskeletal behavior emulation for specificity and diversity. *arXiv preprint arXiv:2603.29332*, 2026. 2, 8

[55] Jie Xu, Viktor Makoviychuk, Yashraj Narang, Fabio Ramos, Wojciech Matusik, Animesh Garg, and Miles Macklin. Accelerated policy learning with parallel differentiable simulation. In *ICLR*, 2022. 2

[56] Pei Xu, Yufei Ye, Shuchun Sun, Yu Ding, Elizabeth Schumann, and C. Karen Liu. MUSIC: Learning muscle-driven dexterous hand control. *ACM Transactions on Graphics*, 45 (4), 2026. 2

[57] Lizhi Yang, Fengze Xie, Yuyang Wang, Lu Gan, Yisong Yue, and Aaron D. Ames. MorphMimic: Morphology-conditioned world model and policy sharing across humanoids. Preprint, 2026. 2

[58] Heyuan Yao, Zhenhua Song, Baoquan Chen, and Libin Liu. ControlVAE: Model-based learning of generative controllers for physics-based characters. *ACM Transactions on Graphics*, 41(6), 2022. 2, 3

[59] Heyuan Yao, Zhenhua Song, Yuyang Zhou, Tenglong Ao, Baoquan Chen, and Libin Liu. MoConVQ: Unified physics-based motion control via scalable discrete representations. *ACM Transactions on Graphics*, 43(4), 2024. 2

[60] Felix 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. 1

VolMusCo: Learning Full-Body Motion Control with Volumetric Muscles

Supplementary Material

A. Muscle Model Construction and Mechanics

We use anatomical volumes to size the muscle bellies. Bone-mounted attachments, routed tendons, and wrapping cylinders transmit their forces to the skeleton.

A.1. Volumetric Asset Construction

A.1.1. MASS284-to-RLMuscleV2 Asset Construction

We start with the 284 muscle paths in MASS284[21] and keep the original skeleton, bone waypoints, and actuator order. We turn each path into a tetrahedral tube and use the RLMuscleV2[26] anatomical meshes to set its size (Fig. S2). VIPER [S1] also samples cross-sections along a muscle, connecting points across slices to form a bundle of elastic rods.

Building a tube along the path. We sample each path at equal distances and place a cross-section at each sample (Fig. S1). At sample i, $\mathbf{t}_i$ points along the path, and $\mathbf{n}_i$ and $\mathbf{b}_i$ are perpendicular directions in the cross-section plane. As the path turns, we apply the smallest rotation R_i that maps $\mathbf{t}_i$ to $\mathbf{t}_{i+1}$. This gives $\mathbf{n}_{i+1} = R_i\mathbf{n}_i$ and $\mathbf{b}_{i+1} = \mathbf{t}_{i+1} \times \mathbf{n}_{i+1}$, keeping successive cross-sections aligned through bends.

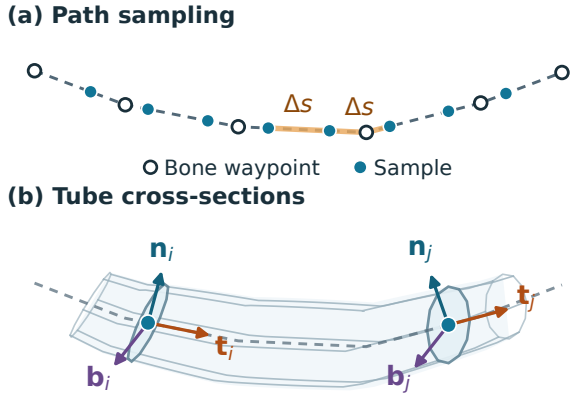

Figure S1. **Sampling the MASS284 path and orienting tube cross-sections.** (a) The dashed line joins the original bone waypoints (open circles). We place samples (filled circles) every Δs of distance measured along the polyline. (b) The blue cross-sections are drawn at samples i and j. At each sample, $\mathbf{t}$ points along the path, and $\mathbf{n}$ and $\mathbf{b}$ span the cross-section.

The cross-sections are narrow at the ends and widest in the middle. Let s be the distance along the path from its starting point and L the total path length. The radius at s is

$$r(s) = \tfrac{1}{2}\rho L\left[0.25 + 0.75\sin^{0.8}(\pi s/L)\right]. \qquad \text{(S1)}$$

Here $\rho = 2r(L/2)/L$ is the midpoint diameter divided by path length. Each calibration tube has ten 12-sided cross-sections and two end tips, giving 122 vertices and 290 tetrahedra per unit.

A separate high-resolution fusiform model with explicit tendon material assigns active contraction to the central 64% of the rest path, with a passive tendon spanning 18% of the path at each end.

Setting the muscle volumes. We assign each MASS284 path to an anatomical muscle group in RLMuscleV2's `humanV2` asset [26] by voting over the nearest mesh vertices along the path. We measure the group's reference volume V_G by summing its tetrahedra in the rest pose. When several MASS284 units map to the same group, we split V_G in proportion to their nominal Hill strengths. Here m and j index units assigned to group G, and $f_{0,m}$ is the nominal strength of unit m. The initial share is V_m^{raw}. We apply a common volume scale and copy the targets of paired left units to the right to obtain the final target V_m for each unit. The volume share and effective cross-sectional area are

$$\begin{aligned} V_m^{\text{raw}} &= V_G \frac{f_{0,m}}{\sum_{j\in G} f_{0,j}}, \\ A_{\text{eff},m} &= \frac{V_m}{L_{f,m}}. \end{aligned} \qquad \text{(S2)}$$

We estimate the fiber span $L_{f,m}$ from the matched anatomical group by projecting its vertices onto the mean fiber direction and taking the span between the 5th and 95th percentiles, with a 2 cm minimum.

We use $A_{\text{eff},m}$ only to initialize the radius of the fusiform calibration mesh. A regular 12-sided cross-section with circumradius r has area $3r^2$. Setting the midpoint cross-sectional area to $A_{\text{eff},m}$ and using $r(L/2) = \rho L/2$ gives the initial radius scale

$$\rho_0 = \frac{2}{L}\sqrt{\frac{A_{\text{eff},m}}{3}}.$$

We clamp ρ_0 to $[0.02, 0.35]$ before adjusting ρ to match the target rest volume V_m.

For the 9-tet mesh, the belly width is $w_m = \sqrt{3V_m/(2L_m)}$, where L_m is the straight-line distance between the two belly tips. In the elongated 9-tet variant, we use 64% of the total rest polyline length as the default target for L_m when selecting the tip locations.

The 284 calibration tubes have a combined rest volume of 30.24 L, summed over their muscle tetrahedra. The anatomical volume calibration excludes the organ group. For simulation, we replace each tube with a belly made of two quadrilateral cross-sections and two tips, with fibers along the belly. We scale the coarse belly to the corresponding tube's volume, so the two mesh sets have the same com-

Table S1. Full-body MASS284 meshes for geometry calibration and simulation.

Mesh	Units	Vertices	Tetrahedra
290-tet tube	284	34,648	82,360
9-tet diamond	284	2,840	2,556

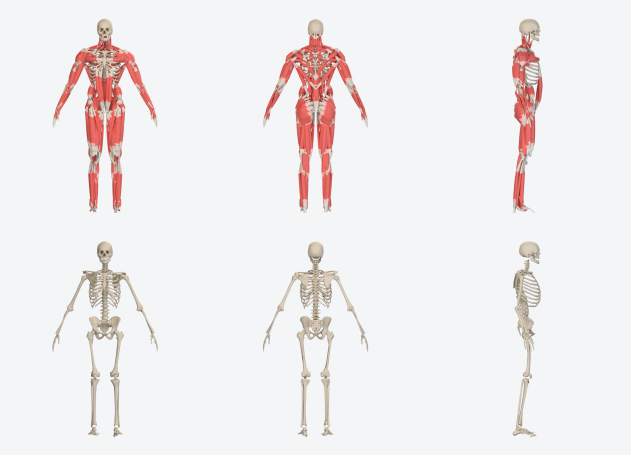

Figure S2. Anatomical calibration of the MASS284 muscle geometry.

bined rest volume. The outer parts of the original paths remain as tendon waypoints. Selected shoulder and hip bellies are curved or flattened to fit their routes and the nearby bones. We rescale their cross sections to match the target volume and check that all rest tetrahedra have positive signed volume. Table S1 gives the full-body mesh counts. Figure S8 compares the three discretizations on one muscle.

Each endpoint uses its configured routed tendon or direct attachment. For a tetrahedron e in muscle unit m, the material and active weights are $w_e = V_e^0 Y_m/(1+\nu)$ and $c_e = V_e^0 S_m$. Here V_e^0 is the rest volume, Y_m is the fitted Young's modulus, and S_m is the active-fiber coefficient. The relative weight of the volume-preserving term is $\nu = 0.3$. We use the MASS284 Hill-type line-muscle model as the calibration reference. At sampled skeletal poses and activation levels, we evaluate each line muscle's force at zero contraction velocity and compute its joint torques from its moment arms. We solve the volumetric muscles at the same poses and activations and fit Y_m, S_m, attachment stiffness, and routed-tendon slack lengths to match the reference forces and torques. We select the fitted parameters on held-out poses.

A.1.2. Surface Geometry

For visualization, we merge the bellies of units that belong to the same anatomical muscle into a continuous surface. Let $f_m(x)$ be the implicit surface field of tube m at position x. We blend these fields with a smooth minimum using $k = 100\,\mathrm{m}^{-1}$.

$$f_{\text{union}}(x) = -\frac{1}{k}\ln\sum_m \exp[-k f_m(x)]. \tag{S3}$$

The blend weight is zero over the outer 20% of the path at each end and increases smoothly over the next 8% of the path length, keeping the tendon ends separate. This step affects only the rendered surface. The simulation uses the original tetrahedral mesh and contact constraints.

A.2. Tendon Routing and Direct Attachments

For one routed endpoint, write its chain as $p_0, \ldots, p_n$, beginning at the muscle tip and continuing through bone-fixed waypoints. Segment directions and tension are

$$d_j = \frac{p_{j+1} - p_j}{\|p_{j+1} - p_j\|}, \qquad T = k_s[L - L_0]_+. \tag{S4}$$

Here j indexes chain segments, d_j is their unit direction, L is total routed length, L_0 is slack length, and k_s is tendon stiffness. The positive-part operator gives zero tension for a slack route. Differentiating its energy gives

$$f_0 = T d_0, \quad f_j = T(d_j - d_{j-1}), \quad f_n = -T d_{n-1}. \tag{S5}$$

Here f_0 acts on the tissue tip, f_j acts on interior bone waypoint j for $1 \le j < n$, and f_n acts on the terminal bone attachment. Each interior waypoint receives the vector sum of the pulls from its two adjacent segments. A direct attachment is the single-segment case ($n = 1$) with zero slack length ($L_0 = 0$).

A.3. Cylinder Wrapping and Contact Reactions

Each muscle with wrapping checks 20 outer belly edges and up to two segments connecting a muscle tip to its first bone waypoint (Fig. S3(a,b)). For one segment and cylinder,

$$\begin{aligned} p(\eta) &= (1-\eta)p_0 + \eta p_1, \qquad 0 \le \eta \le 1, \\ \rho &= \min_{\eta\in[0,1]} \|(I - d_c d_c^\top)(p(\eta) - c)\|, \\ \delta &= [R - \rho]_+. \end{aligned} \tag{S6}$$

Here p_0, p_1 are segment endpoints, η is the interpolation coordinate, d_c is the unit cylinder axis, c is a point on that axis, and I is the identity matrix. The projector $I - d_c d_c^\top$ removes axial displacement, ρ is the closest radial distance, R is effective cylinder radius including clearance, and δ is penetration depth.

We split the contact force at $p(\eta^*)$ between the two endpoints (Fig. S3(c)) using

$$f_0 = (1-\eta^*)k_c\delta n, \qquad f_1 = \eta^* k_c \delta n. \tag{S7}$$

Here η^* is the minimizing segment coordinate, n is the outward radial unit normal, k_c is penalty stiffness, and f_0, f_1 are endpoint contact forces. If an endpoint belongs to a

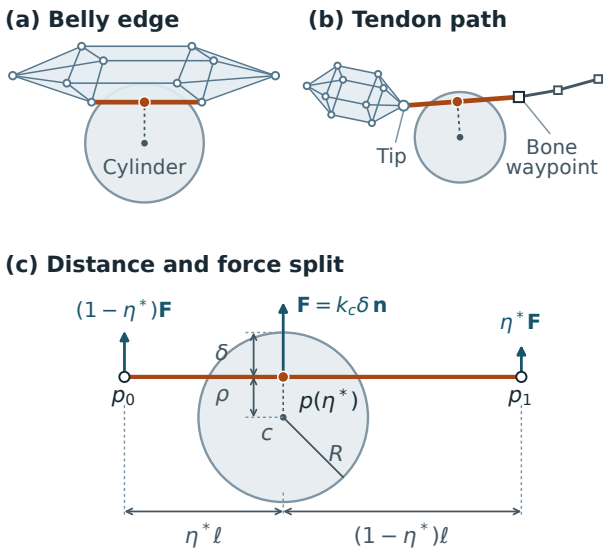

Figure S3. **Cylinder wrapping and contact reactions.** Views along the cylinder axis. The highlighted segment is a muscle-mesh edge in (a) and the tip-to-waypoint span of a routed tendon in (b). In (c), the closest point $p(\eta^*)$ sets the penetration δ and the force weights on p_0 and p_1. The segment length is $\ell = \|p_1 - p_0\|$.

Table S2. Endpoint and cylinder-contact assignments in the wrapping-enabled model asset used for training. Fig. 3 in the main text illustrates the three configurations. Each unit is one muscle belly with an origin and an insertion. The tibialis anterior has a direct origin and a routed insertion. Deltoid and gluteus minimus each have three wrapped units per side.

Assignment	Ends	Units	Representative muscles
Direct spring	178	14	Piriformis
Routed tendon	390	270	Biceps brachii, soleus
Cylinder wrap	–	12	Deltoid, gluteus minimus

bone, its contribution is applied to that bone. The cylinder's owning bone receives the opposite force and torque. Degenerate radial normals require a configured fallback direction. This soft segment test permits penetration and does not guarantee collision-free triangle interiors or fascia interaction.

Model assignments. Table S2 counts individual ends by connection type and assigns each muscle unit to one endpoint category. A unit counts as direct if both ends attach directly to bone and as routed if at least one end uses a routed tendon. The cylinder wrapping row counts wrapped units independently of their endpoint connections.

B. Simulation Engine and GPU Execution

At each tissue update, we batch the coarse tetrahedral muscle meshes on the GPU and solve for quasi-static equilibrium at the current skeletal pose. We use Projective Dynamics to compute muscle deformation and run an additional contact solve for muscles that wrap around cylinders. The rigid-body integrator then advances the skeleton over several substeps using the resulting muscle loads. The muscle solve omits tissue inertia and gravity, while rigid-body integration accounts for skeletal inertia and gravity.

B.1. Runtime Muscle Meshes and GPU Batching

The simulator solves for quasi-static muscle equilibrium at the current skeletal pose and holds the resulting loads across rigid-body substeps. The muscle solve omits tissue inertia and gravity, while the skeletal dynamics retain both.

The simulator solves for the shape of all muscles at every tissue update during rollout collection. We use a coarse mesh to reduce the number of element projections and vertex displacements per iteration. Each belly has ten vertices and nine tetrahedra, arranged as two quadrilateral cross sections and two tips. This reduces the tetrahedron count per environment from 82,360 in the calibration mesh to 2,556 at runtime (Table S1). With the same block size for all muscles, the GPU can batch the solves across muscles and environments. The coarse mesh retains the belly span, target volume, reference fiber directions, and bone-local tendon waypoints.

B.2. Projective Dynamics and Wrapping Solves

Algorithm 2 describes one simulator control step. We read the skeletal pose and velocity and hold the policy's muscle activations fixed for that step. At each tissue update, we refresh the attachments, tendon waypoints, and wrapping cylinders from the current pose. We first run the muscle solve for all units in parallel, initializing each unit m with x_m^-, the most recently accepted 3D positions of its ten mesh vertices. We then run a wrapping solve for the units with configured cylinders C_m.

With the skeleton and activation fixed, the equilibrium solve moves the muscle vertices until the net force on each vertex is small. We use Projective Dynamics (PD) [1] because each iteration can update the nonlinear material response locally while reusing the base Cholesky factorization for the vertex solve. For each tetrahedron, we compute its deformation gradient and volume ratio, then find a local rotation and a volume-preserving projection. We combine the resulting material forces with the forces from the active fiber energy (Eq. 8), tendons, attachments, and the volume barrier to form the nodal force residual. In the muscle solve, we use a 10×10 system with the same factorization for all three coordinate components, using rank-one endpoint corrections for routed tendons.

The wrapping solve adds cylinder-contact energy. After a reset, the muscle solve provides the initial candidate. On subsequent updates, we start from the last accepted wrap-

ping shape moved with the tendon leads. A contact-aware initialization can adjust either candidate before the wrapping iterations. We use the accepted wrapping shape to compute the loads of those muscles. Backtracking uses up to 40 trials per iteration, halving the step size after each rejection. A trial must pass the Armijo test with coefficient 10^{-4} and satisfy $J_e = \det F_e > J_{\min}$ in every tetrahedron. The wrapping solve reuses determinant polynomials to check the volume bound along the search direction. Iteration stops at the residual tolerance or iteration cap.

B.3. Multirate Muscle–Rigid Body Coupling

The policy supplies activation at 50 Hz. At each scheduled tissue update (100 Hz), the solver refreshes tendon waypoints, attachment targets, and wrapping geometry from the current skeletal pose, solves for quasi-static muscle equilibrium, and maps the resulting reactions to a generalized muscle load. The fixed-hold configuration reuses this load for ten rigid-body substeps at 1000 Hz in the Warp implementation of MuJoCo [S7], so the tissue solve and the force mapping run only at tissue updates. The backend also supports an earlier refresh when pose changes exceed a configured threshold. With fixed hold enabled, the LaFAN1 tracking run achieved approximately 800 physical control transitions per second across 2,048 parallel environments on one NVIDIA GeForce RTX 4090. This end-to-end rate includes the time spent updating the world model and actor–critic, which reduces the number of simulation transitions collected per second. The run used 284 muscle units and 12 configured cylinder-wrapping pairs per environment. The nodal residual target was 1 N, with iteration caps of 600 for the muscle solve and 50 for the wrapping solve. The wrapping configuration also accepted finite iterates above the volume threshold $J_{\min} = 0.1$ when the residual target was unmet. In this window, 95.1% of wrapping blocks met the residual target.

C. Reference Motion Processing

Motion-capture clips and video-derived poses provide the tracking references. We convert both sources into joint trajectories for the tracking tasks.

C.1. LaFAN1 References and Motion Conversion

We train on motion clips from LaFAN1 [15]. For evaluation, we retain 47 clips from this training set after excluding `ground`, `fallAndGetUp`, and `pushAndFall` clips.

Before resampling the reference motions, we unwrap each joint-angle channel by taking the shortest signed angular increment between consecutive frames.

$$q^u_{t+1} = q^u_t + \mathrm{wrap}_{(-\pi,\pi]}(q_{t+1} - q_t). \tag{S8}$$

We linearly interpolate the unwrapped angles at the control rate and compute joint velocities from that continuous sequence. This avoids spurious velocity spikes when an angle crosses the $\pm\pi$ boundary.

C.2. AMASS and Motion-X References

We use flat-ground motion clips from the CMU, SFU, and TotalCapture subsets of AMASS[S4]. We train on 1,758 clips and hold out 159 clips for validation and 283 for testing.

We combine 1,874 indexed clips with 326 additional clips from local recordings. For the indexed clips, we use captions from Motion-X[S3] and HumanML3D[S2]. For the additional clips, we use the original trial descriptions and take names. We retain clips that pass pose, ground, and kinematic checks. The full set contains 2,200 clips and 555,108 frames at 30 Hz, totaling 5.14 hours.

C.3. Video Pose Estimation and Retargeting

We extract 2D image landmarks and pelvis-relative 3D body landmarks using the body-pose estimator from BlazePose GHUM Holistic [12], implemented as MediaPipe Pose Full. We adapt the landmarks to the bone lengths and recover joint angles with joint-constrained inverse kinematics, followed by temporal smoothing. For walking clips, we estimate root translation from stance-foot motion. We resample the trajectories at 50 Hz and convert them to the policy skeleton, preserving root poses and checking joint limits and forward-kinematic consistency.

For strength training motions, direct frame-by-frame pose estimation from video can produce noisy trajectories. We therefore select representative poses from demonstration videos, adapt these key poses to the target skeleton, and smoothly interpolate between them to construct continuous motion references for exercises such as dumbbell lateral raises and deadlifts.

We also retarget Reallusion sword-and-shield clips distributed with ASE [S5] and backhand motion from Tennis-MoCap [S6] to the target skeleton. The tennis reference is resampled to 50 Hz and fitted with joint limits and two-hand racket-grip constraints.

D. Policy and World-Model Training

We fit the world model to transitions collected in the coupled simulator. The policy and critic learn from its imagined rollouts using rewards defined by the reference motion. The same training pipeline also supports task-specific goals and rewards.

D.1. Training Pipeline

MuscleVAE [7] supplies the learned-dynamics control precedent, and we use a GPU-accelerated implementation of SuperMimic [27] for shared replay, actor, world-model, and critic updates. We use this learning procedure for volu-

metric muscle control without the generative VAE component.

Algorithm 1 summarizes the training loop. In each iteration, we run the policy in the coupled simulator to collect transitions, fit the model to these transitions, and then update the actor and critic on trajectories predicted by the fitted model. The stepping rates are given in Section 4.2, and the network sizes and training settings are listed in Table S3.

Table S3. **Volumetric tracking hyperparameters.** Settings for the volumetric policy. Network widths list hidden layers.

Setting	Value
Policy encoder widths	$2048, 1024, 512, 256$
Policy decoder widths	$2048, 1024, 512, 256$
Goal latent / action size	$32/284$
World-model widths	$1024, 1024, 1024$
Critic widths	$256, 128$
Parallel environments	2,048
Replay per environment	500 transitions
Total replay capacity	1,024,000
Policy / world-model batch	2,048 each
World-model training window	4 control steps
Policy / critic horizon	16 / 32 control steps
World-model batches / iteration	5
World-model-only warm-up	200 iterations
Base learning rate	4×10^{-4}
Policy / model optimizer	Muon + Adam
Critic optimizer	Adam
Discount γ / return λ	$0.98/0.90$
Policy EMA decay	0.995
Exploration at $a = 0.5$	Beta, $\sigma = 0.0707$

D.2. World-Model Fitting and Policy Optimization

We fit the world model by unrolling it over recorded activation sequences from replay and comparing its predictions with the simulator states, using a four-step window. Given the current state and the current and next muscle activations, the model predicts the next root position and orientation, root linear and angular velocities, joint angles and velocities, and normal ground-contact forces on the foot bodies. In the direct-activation setting, the next activations are the 284 excitations clipped to $[0, 1]$. Muscle deformation and forces are not predicted by the model. They are computed only by the physical solver during data collection.

Each replay transition also stores two flags. The `done` flag stops the unroll at the end of an episode, and the `valid` flag excludes rejected muscle solves and non-finite transitions from world-model fitting.

The policy takes the current physical observation and the previous excitation as input, without stacking past frames. Its goal consists of the motion phase, the current reference frame, and four future reference frames spaced $0.2\,\mathrm{s}$ apart.

To update the actor, we freeze the world model and backpropagate the imagined return in Eq. (4) through its predicted states and its action inputs, so the gradients pass through the model Jacobians instead of stopping at each predicted state. Each imagined reward is weighted by the product $\prod_{j<k} c_j$ of the preceding soft continuation probabilities $c_j = 1 - d_j$, where d_j combines the soft termination terms of Appendix D.3. At every imagined step, we also clip the gradient carried to the next step to a norm of 100 per trajectory, so that a few trajectories cannot dominate an update. Appendix E.2.4 discusses how the horizon affects these gradients.

The critic is trained on the same imagined rollouts. Its targets are finite-horizon λ-returns [43]:

$$\begin{aligned} G_H^\lambda &= V_\psi(\hat{s}_H, g_H), \\ G_k^\lambda &= r_k + \gamma c_k \left[(1-\lambda) V_\psi(\hat{s}_{k+1}, g_{k+1}) + \lambda G_{k+1}^\lambda\right], \end{aligned} \tag{S9}$$

computed backward from $k = H-1$ to 0. Bootstrapped values in both the actor objective and these targets are clamped to the largest attainable return, $r_{\max}/(1-\gamma)$, where $r_{\max}$ is the maximum per-step reward. Table S3 lists the horizon, γ, and λ.

D.3. Tracking Objective

Local articulation is compared in heading frames, while global root trajectory remains explicit. Let p and H be simulated root position and yaw rotation, and let barred symbols denote reference quantities. For body i, define

$$\begin{aligned} z_i^p &= H^\top (p_i - p), \quad \bar{z}_i^p = \bar{H}^\top (\bar{p}_i - \bar{p}), \\ e_i^v &= H^\top v_i - \bar{H}^\top \bar{v}_i, \\ e_i^\omega &= H^\top \omega_i - \bar{H}^\top \bar{\omega}_i. \end{aligned} \tag{S10}$$

Orientations are expressed in the same heading-aligned frame. Each velocity must be rotated by its own heading: if we rotated both the reference and simulated velocities by the simulated heading alone, any yaw offset between them would persist as a world-frame mismatch. The rotated velocities are instantaneous vectors expressed in this frame, not time derivatives of root-relative positions. As a result, rigidly rotating the entire reference about the vertical axis leaves all local terms unchanged. Only the root-heading term changes, and it captures the relative yaw between the reference and the simulated character.

The objective uses exponential body-error kernels and separate root terms. With horizontal root distance d,

Table S4. Shared tracking terms in the endpoint configuration. The continuous weighted sum is multiplied by the control interval Δt.

Component	Weight	Scale / form
Body position	1.0	$\sigma = 0.3\,\mathrm{m}$
Body orientation	1.0	$\sigma = 0.4\,\mathrm{rad}$
Local linear velocity	0.25	$\sigma = 1.0\,\mathrm{m/s}$
Local angular velocity	0.25	$\sigma = 3.14\,\mathrm{rad/s}$
Root position	0.5	Tent radius $4\,\mathrm{m}$
Root yaw	1.0	$\sigma_\psi = 0.25\,\mathrm{rad}$
Hand/head endpoints	1.5	$\kappa = 1000\,\mathrm{m}^{-2}$
Action rate	0.03	$-\|a_t - a_{t-1}\|_2^2$

wrapped yaw error $\delta\psi$, and endpoint set $\mathcal{E}$,

$$\begin{aligned} K_\sigma(d) &= \exp(-d^2/\sigma^2), \\ r_{\text{root}} &= \max(0, 1 - d/D), \\ r_{\text{yaw}} &= \exp[-\rho_\delta(\delta\psi/\sigma_\psi)], \\ r_{\text{tip}} &= |\mathcal{E}|^{-1} \sum_{i\in\mathcal{E}} \exp(-\kappa\|z_i^p - \bar{z}_i^p\|^2). \end{aligned} \tag{S11}$$

The root tent has constant radial slope inside its support and zero slope outside. For the yaw kernel, let $z = \delta\psi/\sigma_\psi$ and $c = 2$. Then $\rho_\delta(z) = z^2$ for $|z| \le c$ and $\rho_\delta(z) = c^2 + 2c(|z| - c)$ otherwise. The endpoint score averages two fingertip scores and one head-top score.

Optional arm and trunk orientation terms have zero weight in the archived endpoint configuration. The positive tracking weights sum to 5.5, giving a maximum step score of 0.11 at 50 Hz. The configured hard termination thresholds are $0.25\,\mathrm{m}$ for root-height error, 0.8 for the tilt feature, and $10\,\mathrm{m}$ for horizontal drift. The analytic soft versions use widths $0.02\,\mathrm{m}$, 0.04, and $1\,\mathrm{m}$, respectively. Clip end is a timeout. Physical transitions rejected by the solver are flagged in replay.

D.3.1. Task-specific rewards and settings

We use the tracking reward in Table S4 as the default. For better visualization, please refer to our accompanying video. Parentheses indicate the corresponding task's position in the video. Fig. S4 shows sequences from the tennis, three-ball juggling, barbell deadlift, and sword-and-shield videos.

LaFAN1. LaFAN1 uses the default tracking reward without hand loads. ($0'55''$)

Human video. We track the retargeted skeleton for walking and in-place exercise clips. ($2'07''$)

Hand loads. We attach grasped props to the hands and incorporate their mass and inertia into the hand rigid bodies. This setup also applies to kettlebells and rackets. The world model, policy, and critic receive ten load features per hand: mass, three center-of-mass coordinates, and six independent inertia entries in the hand frame. The reference clip selects the same load in simulation and imagined rollouts.

Strength training. Lateral raises use a $3\,\mathrm{kg}$ dumbbell in each hand. Dumbbell deadlifts use 12 or $22\,\mathrm{kg}$ per hand. Each dumbbell is fixed to the hand that holds it. Barbell deadlifts use total loads of 40, 80, 100, and $140\,\mathrm{kg}$. In the $140\,\mathrm{kg}$ trial shown in the video, the character remains flexed at the hips and does not complete the reference lift. ($1'20''$) We incorporate its mass and inertia into the right-hand body and keep its center of mass midway between the two grips. A weld constraint maintains the relative pose of the hands and transfers load between them. Kettlebell lifts use a single 12 or $20\,\mathrm{kg}$ kettlebell held with both hands. ($1'05''$)

AMASS box carrying. The carrying reference comes from the ACCAD subset of AMASS. The character holds a $2\,\mathrm{kg}$ box with both hands. We incorporate the box's mass and inertia into the right-hand body, preserving its center of mass, and weld the hands at the reference grasp. The horizontal root-tracking scale is $1.75\,\mathrm{m}$. ($1'46''$)

Sword and shield. We retarget five Reallusion sword-and-shield clips distributed with ASE [S5]. All five clips use the same hand loads, with a $1.8\,\mathrm{kg}$ sword in the right hand and a $3.0\,\mathrm{kg}$ shield in the left. In the sword-tip tracking variant, the goal includes the reference positions of two points near the blade tip, and the reward adds a unit-weight penalty on their mean position error. The wrist-assisted variant adds three motor axes per wrist, with a torque limit of $12\,\mathrm{N\,m}$ per axis. ($1'58''$)

Three-ball juggling. The juggling policy tracks the body motion of the cascade in Table S6. Each ball has a mass of $0.08\,\mathrm{kg}$. Within $40\,\mathrm{ms}$ of a scheduled catch, the script attaches the ball if its center is within one ball radius of the free receiving hand. ($1'52''$)

Tennis. The policy tracks the reference swing from its first frame with a $0.32\,\mathrm{kg}$ racket fixed to the right hand. For the video, we resample the ball's initial position and velocity so that it reaches the racket's string bed at the intended contact time in the reference swing. The video shows two incoming balls and the corresponding returns. ($1'38''$)

E. Evaluation and Ablations

We measure tracking accuracy and success across motion clips. Controlled mechanical tests quantify how the mesh, solver settings, and tissue properties affect muscle loads and joint motion.



Figure S4. **Volumetric motion sequences.** Each row shows ten chronologically ordered snapshots from one recording; timestamps indicate elapsed video time. Rows show tennis, dance, backflip, three-ball juggling, barbell deadlift, and sword-and-shield motion. Muscle color indicates activation from blue to red.

E.1. Evaluation Protocols

E.1.1. Physical Evaluation Protocol

For the LaFAN1 actuation-model comparison, we evaluate 47 references after excluding the `ground`, `fallAndGetUp`, and `pushAndFall` clips. We initialize one environment per clip at its first reference frame and run the frozen policy deterministically. Both backends use a 20 ms control interval, 1 ms rigid-body steps, and full-body ground contact. The volumetric backend updates tissue forces every 10 ms.

We evaluate each clip for up to 10 s (500 control steps). A rollout is successful if it reaches 10 s or the end of the reference clip, whichever comes first, without triggering a failure due to root height, root tilt, root drift, a rejected solve, or a non-finite state. We retain simulated states up to the first termination for error analysis and exclude the terminating step and any later samples.

We report pelvis-relative body MPJPE by subtracting each trajectory's pelvis position and comparing the resulting offsets along the world axes. Root-local MPJPE expresses each offset in its own pelvis orientation frame before comparison. Both metrics exclude the pelvis itself and pool all valid control states across the 47 clips. We also record world-space MPJPE from the unaligned body posi-

tions. The normalized evaluation reward is

$$\mathcal{R}_{\text{eval}} = \frac{\sum_c \sum_{t \in \mathcal{A}_c} r_{c,t}}{r_{\max} \sum_c |\mathcal{A}_c|}, \tag{S12}$$

where $\mathcal{A}_c$ contains the control steps for which clip c is active at the start of the step, $r_{c,t}$ is the step reward, and $r_{\max}$ is the evaluator's reference reward per step. Evaluation throughput counts batched simulation transitions after the first five warm-up steps and includes policy inference, metric evaluation, and resets.

E.1.2. Feedback and Learning Ablation Protocols

The frozen-response baseline uses a fixed nominal generalized-load map $\widetilde{Q}^{\text{mus}}(q, a)$ evaluated at the current skeletal pose q and muscle activation a. In the closed-loop feedback ablation, we fit one MLP per muscle with two 64-unit SiLU hidden layers. The MLP predicts axial force from the distance between the tendon leads, the two signed tendon radii, and activation. We normalize the three lengths by the rest belly span. We reconstruct endpoint forces and routed tendon tensions using the original connection laws and map the reactions to skeletal loads through the attachment points and waypoints. Both variants solve the tissue equilibrium with doubled Young's modulus. Coupled applies the resulting loads, and Frozen applies the nominal map while using the same policy observations and skeletal model.

We fit the map on nominal-material solves from 187 calibration segments and independent per-muscle sweeps, with 60 separate segments for validation. All 247 segments are disjoint from the evaluation segments. The sweeps sample activation in $[0, 1]$ and use per-muscle bounds on lead distances and routed radii derived from the skeletal kinematic tree with a $2\,\text{mm}$ coordinate margin. The validation criterion is less than 1% relative ℓ_2 error for both attachment forces and joint torques. The nominal joint-torque error is 0.60% on 790 validation poses and 4.76% at the 1,000 held-out rollout initial states.

For the material perturbation, Δ_Q is 100 times the ℓ_2 norm of the difference between the frozen and coupled torques divided by the ℓ_2 norm of the coupled torques. Both norms pool the 50 scalar joint torques across the same 790 matched pose and activation probes, excluding the free root. The error is 11.48% with Young's modulus doubled, with a torque RMSE of $6.63\,\text{N m}$.

The recorded-activation replay in Figure 5 uses a map that includes wrapping. For wrapped muscles, the inputs are the two lead positions in the cylinder frame, the two tendon radii, and activation. The map predicts endpoint forces, tendon tensions, contact forces at the leads, and the cylinder wrench. The wrapping input bounds come from the nominal training poses, padded by half a standard deviation per feature, with activation in $[0, 1]$. Frozen applies the predicted loads directly during this replay. The curve ends when a query leaves the calibration domain, and independent fit validation for this map was not completed.

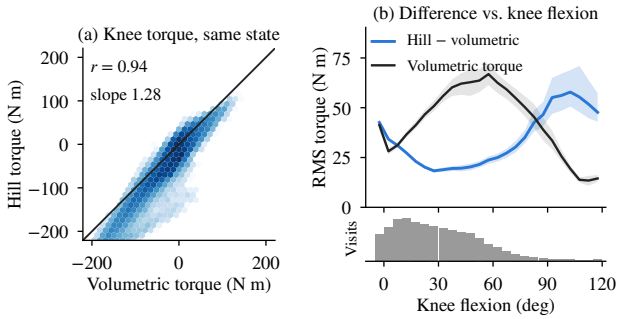

Figure S5. Knee load of the Hill line model versus the volumetric model at the same pose and activation, over states visited by the volumetric policy on LaFAN1 (both knees, 172k samples). (a) Density; line: identity. (b) RMS of the volumetric load (black) and of the Hill-minus-volumetric difference (blue) per 5° of flexion, with 95% intervals from resampling clips; bottom: share of visited states.

We hold policy weights fixed for feedback trials and use matched initial states and paired seeds. Activation/path probes isolate the mechanical response before closed-loop comparisons. Nominal proxy fit and perturbed response errors are reported separately. For learning-efficiency experiments, each algorithm/horizon is trained independently in the same volumetric dynamics. Shared-budget success and time-to-threshold are distinct statistics; failed threshold attainment is retained as censored data. Physical transitions exclude model predictions, and GPU-hours include model fitting and policy/value updates.

E.2. Actuation and Learning Comparisons

E.2.1. Independent Line-Muscle and Volumetric Training

We train a Hill-type line-muscle baseline[60] and a volumetric controller independently with the same model-based learning procedure. Each run uses its own checkpoints and replay. For line-muscle lengths $\ell(q)$, moment arms $R(q)$ and tension f_{Hill}, the baseline force map is

$$Q^{\text{mus}}_{\text{Hill}} = R(q)^\top f_{\text{Hill}}(q, v, a), \qquad \dot{\ell} = -R(q)v. \tag{S13}$$

The two setups use the same skeleton, actuator correspondence, strength calibration, references, rewards, ground-contact settings, control rate, and common policy inputs. For the reported comparison, we use the same initial states and success criteria.

E.2.2. Joint Loads at Matched States

Table 3 of the main paper compares controllers trained independently, each adapted to its own actuation, so it does not

Table S5. Hill versus volumetric generalized load at matched states, right / left. RMS of the volumetric load; Pearson r; least-squares factor s with $Q_{\mathrm{vol}} \approx s\,Q_{\mathrm{Hill}}$; RMS residual after s, relative to the volumetric RMS. Rotations about the model's x, y, z joint axes.

Joint	RMS (N m)	r	s	Residual
Hip x	39.0 / 39.3	0.95 / 0.93	0.79 / 0.71	0.32 / 0.37
Hip z	34.7 / 34.8	0.96 / 0.94	0.76 / 0.71	0.32 / 0.37
Hip y	12.8 / 12.7	0.84 / 0.77	0.47 / 0.41	0.53 / 0.66
Knee	50.0 / 47.7	0.93 / 0.92	0.71 / 0.65	0.36 / 0.41
Ankle x	28.9 / 26.6	0.68 / 0.79	0.50 / 0.64	0.71 / 0.60
Ankle z	8.7 / 10.2	0.24 / 0.34	0.06 / 0.18	0.98 / 0.95
Ankle y	5.0 / 6.2	0.30 / 0.48	0.43 / 0.31	0.96 / 0.88
All leg joints	111.2	0.88	0.65	0.48
Spine and torso	61.1	0.78	0.58	0.62

show how the two models differ at the same state. We run the volumetric policy closed loop on every LaFAN1 clip for up to 60 s, restarting from the reference state after a fall, and record, at the first tissue solve of every other control step, the pose, the activation sent to the muscles, and the generalized muscle load of that solve. At each recorded state we evaluate the Hill load of Eq. S13 with the same activation. Both models use the same skeleton, actuator order, and anchor-to-joint force mapping, so the loads are compared in the same coordinates; the Hill passive load is a median 0.1% of its total.

At the knee (Figure S5) the loads agree in sign and trend ($r = 0.94$), but the Hill load is 28% larger at equal activation, and the difference depends on pose: 18 N m RMS at 25–30° of flexion against a 51 N m load, 43 against 41 N m below 0°, and 48–58 against 13–34 N m beyond 90°. Table S5 lists the lower limbs and trunk. After the least-squares strength factor s ($Q_{\mathrm{vol}} \approx s\,Q_{\mathrm{Hill}}$, per joint), the residual is still 32–41% of the volumetric load at the hip (x) and knee joints and 48% over all leg joints, so the difference is not a constant gain.

E.2.3. Network Architecture Ablation

We compare MLP and Transformer policies for motion tracking with the same 284-actuator wrapping model and world-model configuration. Both runs used 2,048 parallel environments across two GPUs, and 7,000 training iterations. We evaluated both checkpoints from the first frame of the two training clips, a LaFAN1 walking turn and a three-ball juggling motion, for up to 10 s. The juggling score covers the humanoid motion. The MLP achieved lower tracking error on both clips, so we use it in the main experiments.

Both policies passed both clips (Table S6). The MLP reduced overall MPJPE by 14% relative to the Transformer, with the larger reduction on the juggling clip. During the walking turn, mean horizontal pelvis errors were 2.53 m for the MLP and 2.69 m for the Transformer.

Table S6. MLP and Transformer policy comparison on two training clips. MPJPE uses pelvis-relative body positions in world axes and excludes the pelvis. The overall MPJPE pools valid control states across both clips.

Metric	MLP	Transformer
Policy parameters	8,354,620	7,627,676
Training wall time (h)	3.20	3.50
Walking-turn MPJPE (mm)	51.64	54.35
Juggling MPJPE (mm)	10.55	17.99
Normalized return	0.9113	0.9014

E.2.4. Imagination Horizon

The horizon H sets how many steps actor gradients travel through the frozen world model. Let $D_k = \partial \hat{s}_k / \partial \theta$ be the sensitivity of the predicted state to the policy parameters, where the policy map P includes both the actor and the goal construction. Applying the chain rule along the rollout gives

$$D_{k+1} = (f_s + f_u P_s) D_k + f_u P_\theta, \qquad D_0 = 0, \quad \text{(S14)}$$

where f_s and f_u are the Jacobians of the world model with respect to state and action, and P_s and P_θ are the Jacobians of the policy with respect to state and parameters. The recursion for D_k multiplies Jacobians across rollout steps, so errors in the world-model Jacobians can compound over a longer actor horizon H and affect the policy gradient.

E.2.5. Running with Muscle Fatigue

We adopt the 3CC-r fatigue model used in MuscleVAE [7]. Each muscle has active, resting, and fatigued compartments. Their state limits the activation applied to the volumetric solver at each 20 ms control step and carries over across cycles of the 0.80 s LaFAN1 running reference. In a short test from the same initial pose, the frozen policy reached the 20 s limit with fatigue disabled. Scaling both fatigue and recovery rates by 25× and 50× reduced tracking time to 5.72 and 3.56 s, respectively, before root-height termination. Over the shared first 3.56 s, mean local body-position error was 5.34, 5.65, and 9.89 cm for the three settings. All muscle solves were accepted in this short test.

At fatigue and recovery rates, the frozen policy takes shorter strides as fatigue accumulates (Fig. S6). Across the three fatigue stages in Fig. S7, the forward distance between consecutive left-foot touchdowns decreases from 1.62 to 1.10 and 0.78 m.

We fine-tune the pretrained policy, world model, and critic with fatigue enabled during both simulator data collection and imagined rollouts. In a separate paired evaluation at nominal rates with horizontal drift termination disabled, the fine-tuned policy tracks for 124.58 s, compared with 92.60 s for the frozen policy. Both show similar stride shortening. Measuring stride as forward pelvis displacement per reference cycle, the mean decreases from 1.64 to

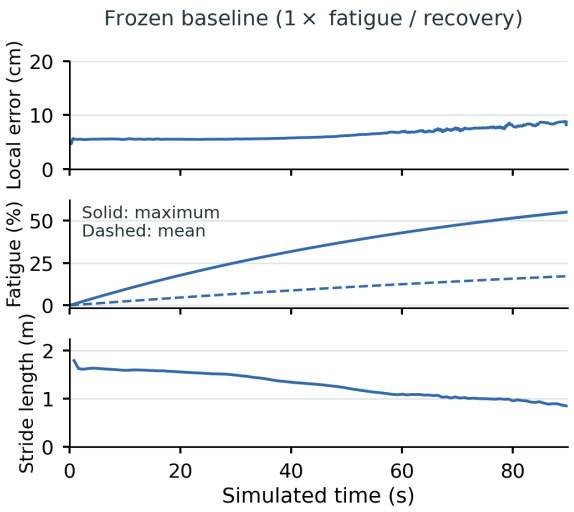

Figure S6. **Long-horizon running with a frozen policy at fatigue and recovery rates ($1\times$).** From top to bottom, the panels show local body-position error smoothed with a centered $0.80\,\mathrm{s}$ moving average, maximum and mean fatigue across muscles, and stride length estimated as forward pelvis displacement per complete $0.80\,\mathrm{s}$ reference cycle. All panels end at $89.82\,\mathrm{s}$, two seconds before the last accepted simulation step. Horizontal drift termination was disabled.

$0.90\,\mathrm{m}$ for the frozen policy and from 1.73 to $1.05\,\mathrm{m}$ for the fine-tuned policy. These means use the first and last ten complete cycles of the shared interval, truncated two seconds before the earlier termination.

E.3. Numerical and Mechanical Ablations

E.3.1. Mesh-Resolution Ablation Protocol

Figure S8 compares the rest shapes of three meshes for the left biceps brachii long head. The 290-tet calibration tube follows the original muscle path and provides the target rest volume. The 9-tet runtime mesh represents the belly with two quadrilateral cross sections and two tips. The 230-tet fit rounds the coarse belly to approximate the calibration tube's profile, preserving the tip positions, rest volume, reference fiber directions, and tendon paths. Table 1 compares the 9-tet and 230-tet simulation meshes, with 2,556 and 65,320 tetrahedra, respectively, across 284 muscle units. Both meshes use the same constitutive parameters, endpoint connections, tendon routing, wrapping cylinders, skeletal inertia, volume barrier, and stepping rates.

In the static comparison, we evaluate both meshes at the rest, walking, and running poses with activation 0.05, and at the rest pose with activation 0.20. At each matched state, T_9 and T_{230} contain the routed tendon tensions. We report the relative discrepancy with respect to the coarse mesh,

$$\Delta_T = 100\,\frac{\|T_{230} - T_9\|_2}{\|T_9\|_2}. \tag{S15}$$

Table S7. **Numerical cost/accuracy ablation.** The reference uses $1000\,\mathrm{Hz}$ muscle updates and a $0.001\,\mathrm{N}$ residual limit. Throughput is the median over three runs in one FP64 CPU environment, measured in accepted $20\,\mathrm{ms}$ control transitions per second.

Muscle rate (Hz)	Residual limit (N)	Δ_T (%)	Work error (%)	Samples/s
50	1	1.89	1.38	29.9
100	1	0.82	0.53	22.2
200	1	0.41	0.20	15.8
100	0.5	0.87	0.56	19.1
100	0.25	0.91	0.61	17.6

The largest discrepancy is 9.46%, at the rest pose with activation 0.20.

For the dynamic comparison in Table 1, we start both meshes from the rest pose, hold activation at 0.05, and simulate free motion for $1\,\mathrm{s}$. We compute the pooled joint-angle RMSE over the 50 completed control steps and 50 scalar joint coordinates, excluding the free root. The RMSE is $1.87°$.

Muscle-driven elbow flexion. We apply the same $50\,\mathrm{Hz}$ activation sequence to the 9-tet and 230-tet models and integrate the elbow independently using the muscle forces computed by each model. Over $6\,\mathrm{s}$, the flexion-angle RMSE between the two trajectories is $2.37°$.

The swept FEM model is built from the 290-tet tube mesh that serves as the basis for our visualizations. The tendons are passive regions of the same mesh. Their end rings follow the bones, and the resulting FEM reaction forces are applied to the skeleton to drive the elbow. After calibration against the 230-tet model, the swept model receives the same activation sequence and yields a flexion-angle RMSE of $9.77°$ relative to the 230-tet trajectory.

E.3.2. Solver Rate and Convergence Ablations

Table S7 varies the muscle-equilibrium rate while holding the prescribed excitation update rate at $50\,\mathrm{Hz}$ and rigid integration at $1\,\mathrm{kHz}$. A second sweep varies the nodal residual limit from 1 to 0.5 and $0.25\,\mathrm{N}$ at a fixed muscle update rate of $100\,\mathrm{Hz}$. All trials use the same material and force laws, waypoint routing, volume barrier, excitation sequence, initial state, and hardware. A rejected solve terminates the trial.

For Δ_T, we evaluate all solver settings along the same reference elbow trajectory with the same prescribed activations. We compare the tension at both tendon ends on the $1\,\mathrm{ms}$ rigid-step grid. We compute the relative ℓ_2 error over time and both ends for each muscle and average over muscles with nonzero reference tension.

In a separate rollout, we let the elbow move freely under the same activation sequence. We compare its angle and net muscle work with the reference. The work is

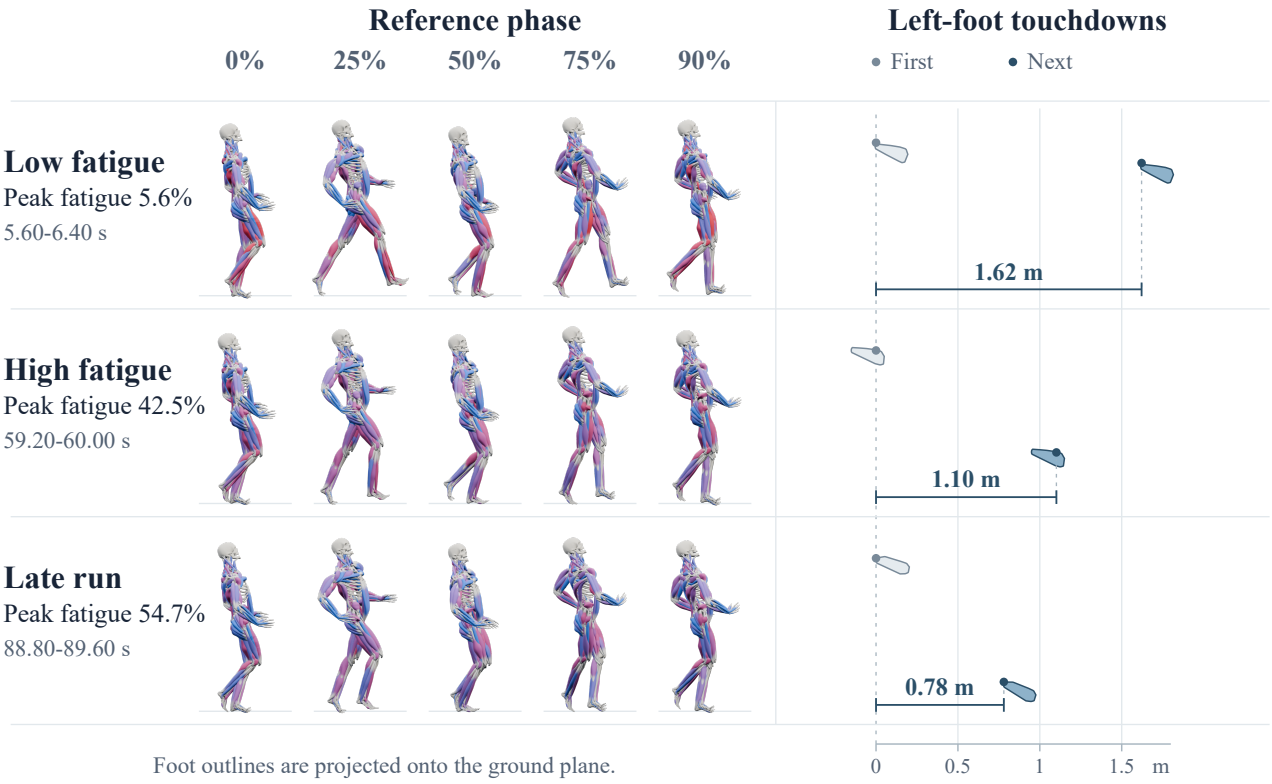

Figure S7. **Running gait across three fatigue stages.** Rows show the three fatigue stages, and columns show snapshots at 0%, 25%, 50%, 75%, and 90% of the reference gait cycle. Peak fatigue is the largest fatigued fraction across the 284 muscle units at the 25% phase. The right panel shows two consecutive left-foot touchdowns on a common spatial scale, with the first touchdown aligned across rows. Each dot marks the point where the foot first contacts the ground at that touchdown. The forward distance between the two dots gives the stride length, which decreases from 1.62 m to 1.10 m and then to 0.78 m.

$W = \sum_i Q_i^{\mathrm{mus}}(q_{i+1} - q_i)$, where Q_i^{mus} is the muscle torque applied during rigid substep i and q_i is the elbow angle in radians. We report $100\,|W - W_{\mathrm{ref}}|/|W_{\mathrm{ref}}|$ as the work error in percent. Each rollout lasts 2 s, and the reference uses a nodal residual limit of 0.001 N. We also record the number of rejected solves and the maximum nodal residual.

Table S7 uses a 1000 Hz reference. For the convergence test, all trials use 100 Hz muscle updates, including the reference, which uses a nodal residual limit of 0.001 N. We compare tendon tensions along the same prescribed elbow trajectory with the same activations. To compare mechanical work, we let the elbow move freely in each run under the same activation sequence and sum the muscle work over that rollout. Relative to the 100 Hz reference, tightening the residual limit from 1 to 0.25 N reduces the tension error from 0.266% to 0.048% and the work error from 0.101% to 0.022%.

E.3.3. Solver Comparison

We benchmarked routed tetrahedral FEM with Newton iterations, the PD solver, and a quasi-static total-Lagrangian MPM implementation. Each world contains two biceps heads. The four inputs combine elbow coordinates of -0.4 and -0.8 rad with activations of 0.1 and 0.5. Runs used an RTX 4070 SUPER with warm and cold starts and batches of one and 256 identical worlds. Timings cover the synchronized solver call with preloaded inputs and exclude routing and post-solve checks. For the 256-world warm-start batch, the median across the four per-input median latencies was 133.0, 2.23, and 113.6 ms for Newton, PD, and MPM, respectively. Each per-input median uses seven repeats.

All three methods met their respective 1 N stopping tests in these four cases. FEM and PD test tetrahedral nodal forces, and MPM tests its independent grid forces. Relative to the Newton solution, the maximum absolute elbow torque discrepancy in these cases was 0.020 N m for PD and 5.09 N m for MPM. Table S8 summarizes the warm-start timings and torque errors. Fig. S9 shows the elbow motion produced by each solver.

E.3.4. Muscle Mass and Inertia

The inertial prototype assigns mass to muscle vertices and adds the backward-Euler inertial term to the tissue objective. With mass matrix M, time step h, and predictor

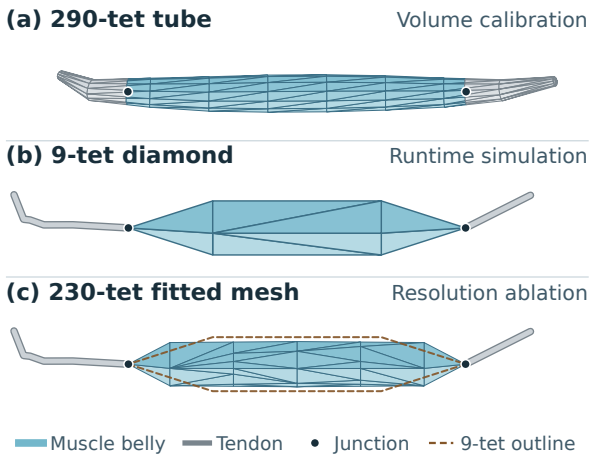

Figure S8. **Three meshes of the left biceps brachii long head.** Blue marks the central 64% of the 290-tet tube's path and gray marks the two tendon regions. Gray lines in panels (b) and (c) trace the same routed tendons from the 9-tet and 230-tet belly tips to the bone attachments. The dashed line in (c) shows the 9-tet outline. Dots mark the two belly–tendon junctions. The views use the same scale and are aligned at the junctions.

Table S8. **Warm-start solver comparison.** Each timed batch contains 256 identical two-biceps-head worlds. Times are medians across four cases, each timed seven times. The stopping threshold is 1 N on tetrahedral nodal forces for Newton and PD and on independent grid-node forces for MPM. Torque differences are maximum absolute elbow discrepancies relative to FEM Newton across the four cases.

Solver	Batch time (ms)	Cases converged	Torque diff. (N m)
FEM Newton	133.0	4/4	Reference
VolMusCo (ours)	2.23	4/4	0.020
MPM	113.6	4/4	5.09

$y = x_n + hv_n + h^2 g$, the term is $(x-y)^\top M(x-y)/(2h^2)$. The updated tissue shape transfers its tendon and contact reactions to the joint dynamics.

In the elbow trial, muscle inertia produced pronounced forearm muscle oscillations that compliant attachments reduced (Fig. S10). In a full-body stress test, the inertial model took 76 s to simulate 1 s at 500 Hz, and the coupled solve failed to converge at the training update rates. We therefore use quasi-static muscle equilibrium for training rollouts while retaining skeletal inertia.

E.3.5. Local Fascial Coupling

We repack the rest meshes of the medial and lateral gastrocnemius heads and the soleus into a single calf compartment. We define a common axis from the mean direction between the tendon tips and align the two cross-sectional rings of all three bellies within their overlapping axial interval. At each ring, three quadrilateral sectors tile a hexagonal outline. We size each sector to preserve the muscle's reference volume. For muscle m, we set the sector area A_m using $V_m = A_m[L_s + (L_m - L_s)/3]$, where L_s is the distance between rings and L_m is the axial tip-to-tip span. The first term accounts for the central prism and the second for the two tapered ends. Each belly retains its ten independent vertices, nine tetrahedra, and original tendon endpoints and routing.

We pair the facing quadrilateral patches of neighboring bellies in the rest mesh. During the solve, we recompute the planes of each patch's two triangles from the deformed vertices and evaluate the signed normal gaps of the opposing vertices. We use the same unilateral penalty for these vertex–plane contacts and for contact between the belly-ring vertices and the sleeve walls:

$$E_{\text{contact}} = \frac{1}{2}\sum_i k_i[-g_i]_+^2, \qquad [z]_+ = \max(z, 0),$$

where g_i is the signed clearance for sample i and k_i is its penalty stiffness. A penetrating vertex receives a normal force of magnitude $k_i[-g_i]_+$. At muscle–muscle interfaces, barycentric weights distribute the opposite force to the triangle vertices. The penalty acts along the normal, so the bellies can slide tangentially and separate.

The hexagonal fascial sleeve encloses the central compartment and follows the tibia. One scalar displacement r allows the sleeve to expand or contract uniformly across its cross section while keeping its axial length fixed. Its tension-only membrane edges resist stretching, and a radial spring restores the rest cross section:

$$E_{\text{sleeve}}(r) = \frac{1}{2}\sum_e k_e[\ell_e(r) - \ell_{e,0}]_+^2 + \frac{1}{2}k_r r^2.$$

Here $\ell_e(r)$ and $\ell_{e,0}$ are the current and rest lengths of edge e, k_e is its stiffness, and $k_r = 2500\,\text{N/m}$. The sleeve limits outward bulging through normal contact and transmits the resulting wall reactions to the tibia.

At each tissue update, we solve for the muscle vertices and sleeve displacement together using the tissue, tendon, contact, and sleeve energies. In the knee–ankle demonstration (Fig. S11), we update the tissue equilibrium at 30 Hz and hold the resulting tendon and sleeve loads for 16 rigid-body steps at 480 Hz.

In a separate benchmark, we replayed 85 knee–ankle poses and muscle activation inputs for the three-muscle calf on one CPU thread in FP64. With the same packed meshes and material tangent correction, the mean warm-start time per tissue solve increased from 19.54 ms for independent bellies to 30.97 ms with sliding contact and the sleeve (1.58×). The coupled solve updates all three bellies together. Fascial coupling could be incorporated into full-body training and evaluation in future work.

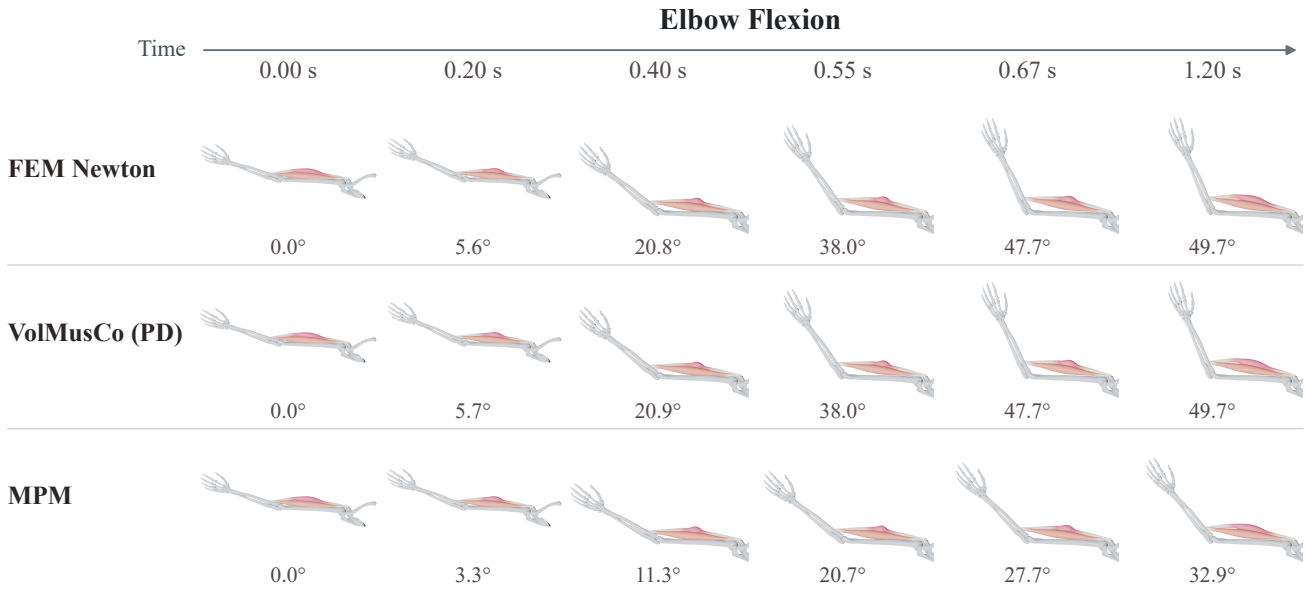

Figure S9. **Muscle-driven elbow flexion across solvers.** From left to right, the columns show elbow flexion at successive simulation times. The rows show FEM Newton, VolMusCo (PD), and MPM. All three solvers receive the same activation and advance the elbow using their own computed muscle forces. PD and FEM Newton produce nearly identical trajectories. MPM flexes more slowly and reaches a smaller flexion angle.

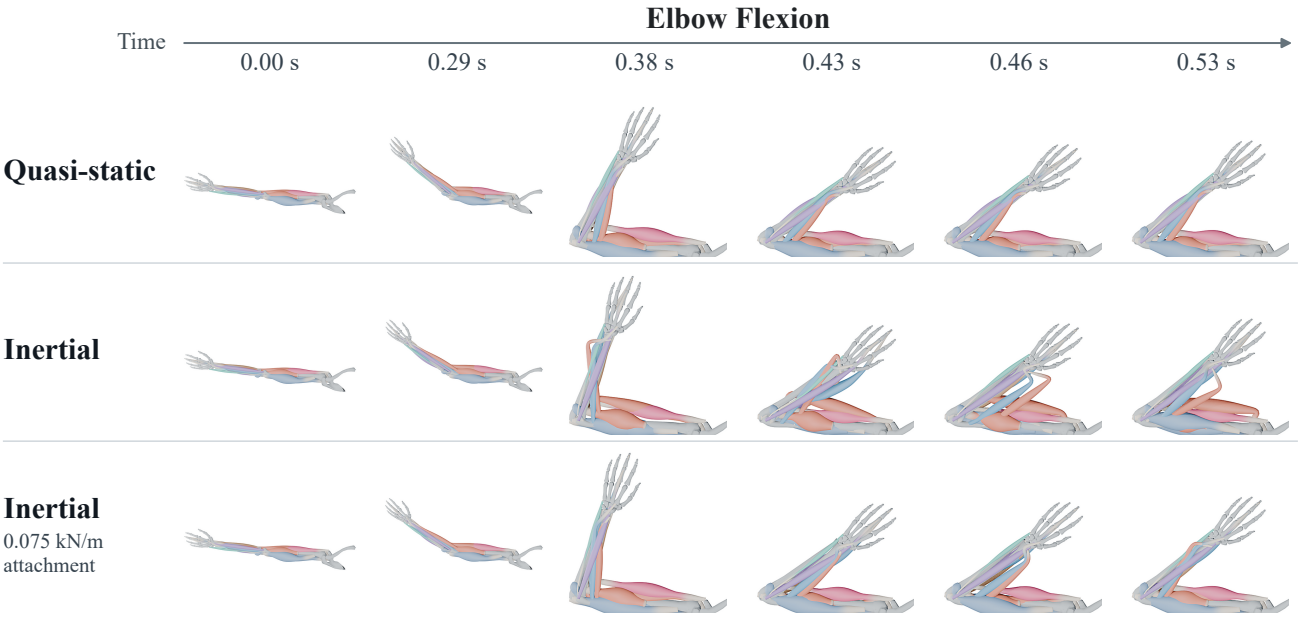

Figure S10. **Effect of muscle inertia during elbow flexion.** From left to right, the columns show elbow flexion at successive simulation times. The rows show the quasi-static model, the inertial model, and the inertial model with a $0.075\,\mathrm{kN/m}$ attachment spring. The compliant attachment reduces forearm muscle oscillations.

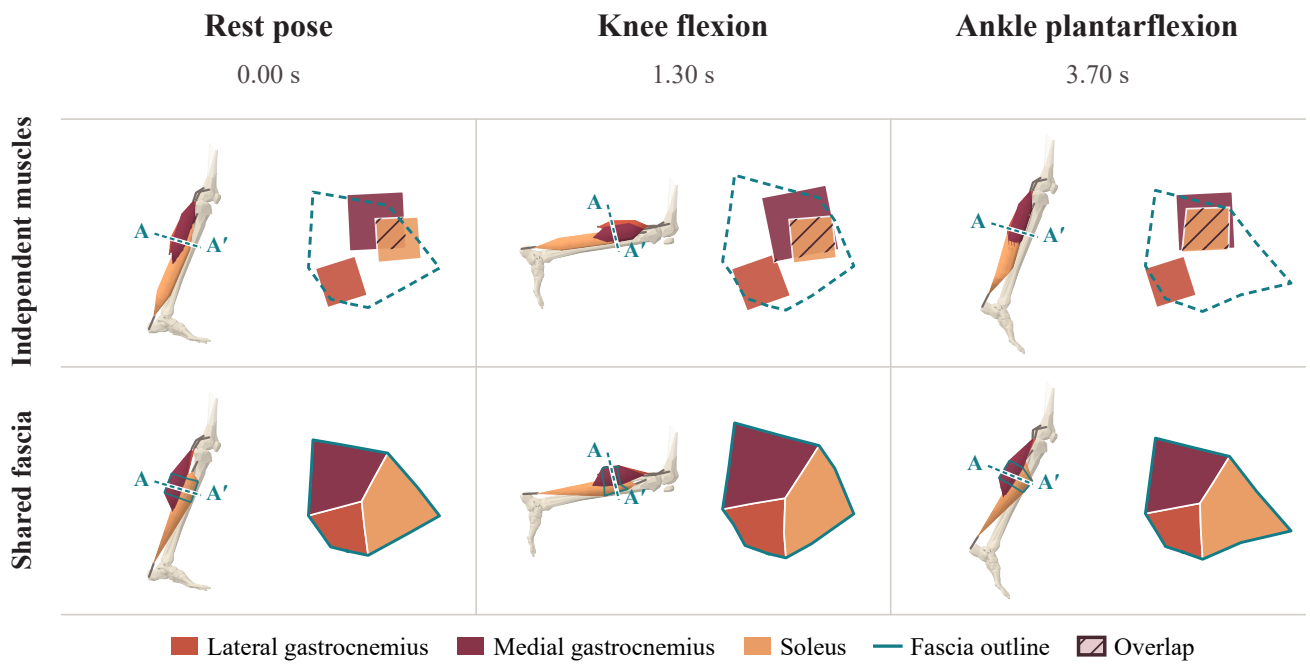

Figure S11. **Calf muscle deformation with and without shared fascia.** From left to right, the columns show the rest pose, knee flexion, and ankle plantarflexion. The top row shows independent muscle bellies, and the bottom row shows muscle bellies in a shared fascial sleeve. In the A–A′ sections, dashed outlines show the corresponding sleeve boundary, and hatching marks overlap between muscle bellies.

Algorithm 1 Train a volumetric-muscle motion-tracking policy

Input: Reference motions $\mathcal{D}$, muscle model $\mathcal{M}$
Output: Policy π_θ, world model W_ϕ, critic V_ψ

function TRAIN($\mathcal{D}, \mathcal{M}$)
  $\mathcal{D} \leftarrow$ RETARGETTOMASS284($\mathcal{D}$)
  Initialize $\pi_\theta, W_\phi, V_\psi$
  RESETTOREFERENCE($\mathcal{D}$)
  $\mathcal{B} \leftarrow$ COLLECTINITIALTRANSITIONS
  for each training iteration **do**
    for each parallel collection step **do**
      $(s, f) \leftarrow$ READTRACKINGSTATE
      $o, g \leftarrow$ OBSERVATIONANDGOAL($s, \mathcal{D}, f$)
      $a \leftarrow$ BETAEXPLORE($\pi_\theta(o, g)$)
      $s', b \leftarrow$ SIMULATE($s, a, \mathcal{M}$)
      $d \leftarrow$ TRACKINGDONE($s', \mathcal{D}, f + 1$)
      $d \leftarrow d \vee \neg b \vee$ CLIPEND($\mathcal{D}, f + 1$)
      APPEND($\mathcal{B}, (s, a, s', f, d, b)$)
      if d **then**
        RESETTOREFERENCE($\mathcal{D}$)
      end if
    end for
    TRAINWORLDMODEL($W_\phi, \mathcal{B}$)
    if world-model warm-up complete **then**
      TRAINACTORCRITIC($\pi_\theta, V_\psi, W_\phi, \mathcal{B}, \mathcal{D}$)
    end if
  end for
  return $\pi_\theta, W_\phi, V_\psi$
end function
function TRAINWORLDMODEL($W_\phi, \mathcal{B}$)
  $(s_0, a_0, s_1, \ldots, a_{K-1}, s_K) \leftarrow$ SAMPLEWINDOW($\mathcal{B}$)
  $\hat{s}_0 \leftarrow s_0$
  for $t = 0, \ldots, K - 1$ **do**
    $\hat{s}_{t+1} \leftarrow$ WORLDSTEP($W_\phi, \hat{s}_t, a_t$)
  end for
  $\ell \leftarrow$ STATECONTACTLOSS($\hat{s}_{1:K}, s_{1:K}$)
  Mask invalid transitions and steps past episode end
  Update ϕ with the masked loss
end function
function TRAINACTORCRITIC($\pi_\theta, V_\psi, W_\phi, \mathcal{B}, \mathcal{D}$)
  $H_\pi \leftarrow 16, \quad H_V \leftarrow 32$
  $(s_0, f_0) \leftarrow$ SAMPLEVALIDSTART($\mathcal{B}$); $\hat{s}_0 \leftarrow s_0$
  Freeze ϕ, ψ for the actor update
  for $t = 0, \ldots, H_\pi - 1$ **do**
    $o_t, g_t \leftarrow$ OBSERVATIONANDGOAL($\hat{s}_t, \mathcal{D}, f_0 + t$)
    $z_t \leftarrow \mathrm{Enc}_\theta(g_t)$
    $a_t \leftarrow \mathrm{sigmoid}(\mathrm{Dec}_\theta(o_t, z_t))$
    $\hat{s}_{t+1} \leftarrow$ WORLDSTEP($W_\phi, \hat{s}_t, a_t$)
    $r_t, \delta_t \leftarrow$ SCORESTEP($\hat{s}_{t+1}, \mathcal{D}[f_0 + t + 1]$)
  end for
  $J_\pi \leftarrow$ SURVIVALRETURN(r, δ, V_ψ, H_π)
  Extend to H_V with π_θ and gradients off
  Backpropagate J_π through H_π steps to update θ
  $G^\lambda \leftarrow$ LAMBDARETURNS(r, δ, V_ψ, H_V)
  Fit V_ψ to G^λ_t for $0 \le t < H_\pi$
end function

Algorithm 2 Step the volumetric musculoskeletal simulator

Input: State s, muscle activations a, muscle model $\mathcal{M}$
Output: Returned state s', step acceptance flag b

function SIMULATE($s, a, \mathcal{M}$)
  $(q, v) \leftarrow$ READSKELETON(s)
  Hold a fixed for the control step
  for each tissue update **do**
    UPDATEANCHORSANDCYLINDERS(q)
    for each muscle m in parallel **do**
      $x_m, b_m \leftarrow$ SOLVE($x_m^-, q, a_m, \mathcal{M}_m$)
    end for
    for wrapping muscles m in parallel, if all b_m are true **do**
      $x_m^0 \leftarrow$ WRAPSEED(x_m, x_m^-, q, C_m)
      $x_m, b_m \leftarrow$ SOLVE($x_m^0, q, a_m, \mathcal{M}_m, C_m$)
    end for
    if any b_m is false **then**
      return READSTATE(q, v, a), **false**
    end if
    $x_m^- \leftarrow x_m$ for all m
    $\{w_i\} \leftarrow$ ENDPOINTCONTACTWRENCHES
    $Q_{\mathrm{muscle}} \leftarrow \sum_i J_i(q)^\mathsf{T} w_i$
    for each rigid substep in this interval **do**
      $q, v \leftarrow$ RIGIDBODYSTEP($q, v, Q_{\mathrm{muscle}}$)
    end for
  end for
  return READSTATE(q, v, a), **true**
end function
function SOLVE($x, q, a, \mathcal{M}, C = \emptyset$)
  while residual above tolerance and budget remains **do**
    $P \leftarrow$ PROJECTROTATIONANDVOLUME(x)
    $F \leftarrow$ NODALFORCES($x, q, a, \mathcal{M}, C$)
    $x \leftarrow$ VOLUMESAFEPDSTEP(x, P, F)
  end while
  return x, ACCEPT($x, q, a, \mathcal{M}, C$)
end function

Supplementary References

[S1] Baptiste Angles, Daniel Rebain, Miles Macklin, Brian Wyvill, Loic Barthe, JP Lewis, Javier Von Der Pahlen, Shahram Izadi, Julien Valentin, Sofien Bouaziz, et al. Viper: Volume invariant position-based elastic rods. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 2(2):1–26, 2019.

[S2] Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating diverse and natural 3d human motions from text. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5142–5151. IEEE, 2022.

[S3] Jing Lin, Ailing Zeng, Shunlin Lu, Yuanhao Cai, Ruimao Zhang, Haoqian Wang, and Lei Zhang. Motion-x: A large-scale 3d expressive whole-body human motion dataset. *Advances in Neural Information Processing Systems*, 36:25268–25280, 2023.

[S4] 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.

[S5] Xue Bin Peng, Yunrong Guo, Lina Halper, Sergey Levine, and Sanja Fidler. Ase: Large-scale reusable adversarial skill embeddings for physically simulated characters. *ACM Transactions On Graphics (TOG)*, 41(4):1–17, 2022.

[S6] Juan Diego Pulgarin-Giraldo, Andres Marino Alvarez-Meza, Luis Gerardo Melo-Betancourt, Santiago Ramos-Bermudez, and German Castellanos-Dominguez. A similarity indicator for differentiating kinematic performance between qualified tennis players. In *Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications*, pages 309–317. Springer, 2017.

[S7] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, pages 5026–5033. IEEE, 2012.