MUSCULOSKELETAL SIMULATION OF LIMB MOVEMENT BIOMECHANICS IN DROSOPHILA MELANOGASTER

Anonymous authors
Paper under double-blind review

ABSTRACT

Computational models are critical to advance our understanding of how neural, biomechanical, and physical systems interact to orchestrate animal behaviors. Despite the availability of near-complete reconstructions of the *Drosophila melanogaster* central nervous system, musculature, and exoskeleton, anatomically and physically grounded models of fly leg muscles are still missing. These models provide an indispensable bridge between motor neuron activity and joint movements. Here, we introduce the first 3D, data-driven musculoskeletal model of *Drosophila* legs, implemented in both OpenSim and MuJoCo simulation environments. Our model incorporates a Hill-type muscle representation based on high-resolution X-ray scans from multiple fixed specimens. We present a pipeline for constructing muscle models using morphological imaging data and for optimizing unknown muscle parameters specific to the fly. We then combine our musculoskeletal models with detailed 3D pose estimation data from behaving flies to achieve muscle-actuated behavioral replay in OpenSim. Simulations of muscle activity across diverse walking and grooming behaviors predict coordinated muscle synergies that can be tested experimentally. Furthermore, by training imitation learning policies in MuJoCo, we test the effect of different passive joint properties on learning speed and find that damping and stiffness facilitate learning. Overall, our model enables the investigation of motor control in an experimentally tractable model organism, providing insights into how biomechanics contribute to generation of complex limb movements. Moreover, our model can be used to control embodied artificial agents to generate naturalistic and compliant locomotion in simulated environments.

1 INTRODUCTION

Understanding how to coordinate multiple limbs with many degrees of freedom (DoFs) to accomplish diverse motor tasks is a long-standing challenge in both motor control neuroscience and robotics. Performing such systems-level investigations hinges upon detailed knowledge of musculoskeletal structures and their mechanical properties. Neuromechanical models, therefore, serve as a crucial tool enabling rapid experimentation and hypothesis testing under controlled conditions (Edwards, 2010; Ausborn et al., 2021; Kim et al., 2022; Delp et al., 2007; Ramdya and Ijspeert, 2023).

Muscle-driven systems are characterized by intrinsic compliance and redundancy which can simplify the control problems faced by artificial agents as they do for animals in the real world (Geijtenbeek et al., 2013; Lee et al., 2018a; Wang et al., 2012). However, many existing models abstract away key elements of musculature and passive biomechanics (Aldarondo et al., 2024; Wang-Chen et al., 2024) limiting their ability to capture the richness of neuromuscular coordination. This, in turn, limits their utility for understanding how the nervous system generates robust, adaptive behaviors. Overcoming this gap will require modeling a system that generates sufficiently complex limb-dependent movements using a tractably small neuromuscular controller.

The fruit fly, *Drosophila melanogaster*, represents an ideal organism for this reverse-engineering challenge. Flies generate complex behaviors using a compact, well-characterized, and genetically

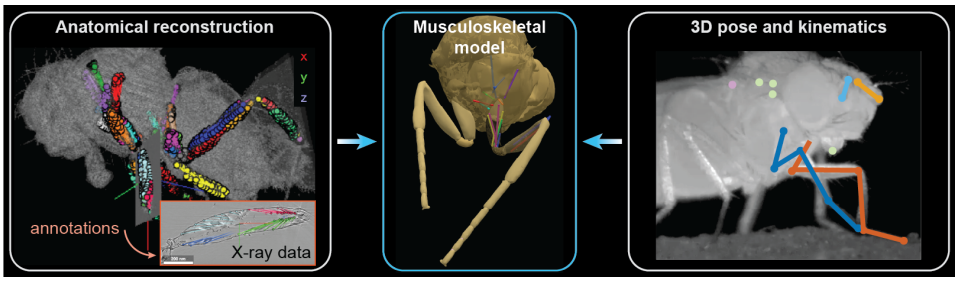

Figure 1: **Pipeline to develop *Drosophila* leg musculoskeletal models.** (Left) Anatomical data from multiple flies were used to determine muscle attachment points and fiber paths, providing constraints for constructing the muscle model in OpenSim. (Right) 3D pose estimation data from behaving animals were then used to train the muscle-actuated agent to recapitulate detailed kinematics in OpenSim and MuJoCo. (Middle) The resulting muscle model informed by both leg anatomy and joint kinematics.

accessible nervous system (Bellen et al., 2010; Rubin and Spradling, 1982). Despite substantial progress in anatomical (e.g., connectomic (Dorkenwald et al., 2024; Phelps et al., 2021)), neural (Aimon et al., 2019; Simpson, 2024; Chen et al., 2018), and behavioral characterization (Aranha and Vasconcelos, 2018; Özdil et al., 2024a), recently published biomechanical models of *Drosophila* legs (Vaxenburg et al., 2025; Wang-Chen et al., 2024; Lobato-Rios et al., 2022) lack the anatomical and physiological accuracy required to model and investigate how motor networks coordinate muscles, limb kinematics, and behavior.

Here, we present the first biologically detailed musculoskeletal model of *Drosophila* legs across two widely used physics engines, OpenSim and MuJoCo. Our model extends NeuroMechFly (Wang-Chen et al., 2024; Lobato-Rios et al., 2022), a whole-body biomechanical simulation of the adult fly, by incorporating anatomically detailed muscle representations and physiological properties. Specifically, we implement a Hill-type muscle model informed by anatomical data (Azevedo et al., 2024; Lesser et al., 2024; Dinges et al., 2021) to actuate seven DoFs in three leg joints. Using this model, we predict the relative contributions of individual muscles to joint movements generated by real behaving flies. Furthermore, we study the influence of passive joint properties on the successful imitation of animal behavior. Our work integrates anatomical, physiological, and behavioral data into a unified modeling framework, providing a critical foundation for future investigations of *Drosophila* neuromuscular coordination and behavioral control. This modeling framework (Figure 1) can also accelerate the development of musculoskeletal models for other species, inform machine learning-based algorithms for embodied agents, and facilitate the control of bioinspired robots.

2 Related Work

2.1 Musculoskeletal Models of Animals

Musculoskeletal models have emerged as powerful tools for investigating animal biomechanics. Open-source simulation platforms such as OpenSim (Delp et al., 2007), MyoSuite (Caggiano et al., 2022), and DART (Lee et al., 2018b), as well as commercial frameworks like HyFyDy (Geijtenbeek, 2021), have made it possible to test state-of-the-art learning algorithms on complex, high-dimensional systems. Use cases include the control of human arms for manipulation and legs for locomotion (Song et al., 2021; Caggiano et al., 2024). Although primarily developed to simulate human bodies (Sylvester et al., 2021; Seth et al., 2018), these tools have enabled the creation of musculoskeletal models for a wide range of species, including rodents (Charles et al., 2018; Tata Ramalingasetty et al., 2021; Gilmer et al., 2025; DeWolf et al., 2024), primates (Ogihara et al., 2009; Chan and Moran, 2006), horses (van Bijlert et al., 2024), and birds like ostriches (Rankin et al., 2016; La Barbera et al., 2021). Researchers have also developed custom tools to study the interplay between neural circuits and biomechanics in worms (Zhao et al., 2024), insects (Guo et al., 2018), and other animals (Arreguit et al., 2023).

Recently, the fruit fly, *Drosophila melanogaster*, has become an increasingly prominent animal model for musculoskeletal simulations. The first anatomically realistic fly body model, NeuroMechFly, used micro-computed tomography (micro-CT) scans to construct a morphologically accurate 3D fly simulation. In its earliest incarnation, simplified antagonistic spring-damper muscles were used in the legs of NeuroMechFly to simulate locomotor dynamics (Lobato-Rios et al., 2022). More recent efforts in this (Wang-Chen et al., 2024) and other (Vaxenburg et al., 2025) fly body models have enhanced simulations by adding, for example, multimodal sensorimotor transformations to accomplish a wider range of behavioral tasks. However, these models still rely on abstract —position- or torque-based— joint controllers and lack realistic muscle actuation. We address this gap in biological fidelity in our work by explicitly modeling individual fly leg muscles based on detailed anatomical data.

2.2 Motor control and muscle synergies

Coordinating the activities of many muscles is a fundamental challenge in motor control. It has been proposed that the nervous system simplifies this complex problem by using muscle synergies—groups of muscles that are co-active, acting as functional units (Bizzi and Cheung, 2013; Ting and McKay, 2007). By combining a limited number of these synergies, the brain can more efficiently produce diverse movements without needing to control each muscle independently (Bizzi and Cheung, 2013; Saito et al., 2018). Muscle synergies have been identified and analyzed in real experimental data (Tresch et al., 2006; Sponberg et al., 2015; Al Borno et al., 2020), and inspired the development of controllers that operate within reduced-dimension muscle activation spaces (Berg et al., 2024). These synergy-based approaches facilitate robust coordinated movements while significantly reducing control complexity. For instance, human walking can be performed across a range of speeds and directions using just a few lower-limb synergies (Saito et al., 2018).

Muscle synergies have also been studied extensively in insects with respect to how sensory feedback reinforces coordinated muscle activation (Bidaye et al., 2018; Zill et al., 2018). By leveraging the complete wiring diagram of the *Drosophila* nervous system—known as the *connectome*—recent studies have identified specific neural circuits that may coordinate fly limb movements (Syed et al., 2024; Lesser et al., 2024). Nevertheless, because of the anatomical and functional complexity of *Drosophila* leg muscles, we cannot predict muscle activity patterns for a given behavior based on neural connectivity or neural recordings alone. In this work, we address this gap by (i) developing an anatomically detailed muscle model of *Drosophila* legs and (ii) simulating this model to reproduce measured joint kinematics across behaviors, allowing us to predict underlying active muscle synergy groupings.

2.3 *Drosophila* leg musculature

Each fly leg is a multi-jointed appendage with at least seven DoFs across five joints (Lobato-Rios et al., 2022; Haustein et al., 2024). These joints are actuated by approximately 19 muscles, which in turn are controlled by approximately 69 motor neurons (Azevedo et al., 2024; Lesser et al., 2024; Soler et al., 2004). Notably, the muscle structure across legs is nearly identical, with a few exceptions including tergal depressor of the trochanter (TDT) muscles in the middle legs that facilitate jump escape (Swank, 2012). Leg biomechanics are further complicated by biarticular muscles that span multiple joints to coordinate movements across body segments (Lesser et al., 2024). Additionally, the passive properties of joints, including their elasticity and damping, can either resist, or assist motion depending on the direction of movement (Ache and Matheson, 2013).

3 Methods

3.1 Acquisition of anatomical data

To reconstruct the front leg muscles, we used two publicly available datasets (Kuan et al., 2020; Dinges et al., 2021) and one custom dataset collected for this study. The custom dataset was acquired using using synchrotron radiation μCT. Each of the three datasets captured different foreleg postures, allowing us to better understand how muscle attachment points vary across joint configurations. We focused on annotating muscle fibers that span the thorax–coxa, coxa–trochanter, and femur–tibia

joints (Figure 2A) and grouped them by function as described in (Azevedo et al., 2024; Soler et al., 2004). For muscles located in the thorax, we used data from (Dinges et al., 2021); for the remaining leg muscles, we relied on (Kuan et al., 2020). Muscle attachment points were then cross-validated using our custom scan. Together, these datasets provided the anatomical basis for our muscle model.

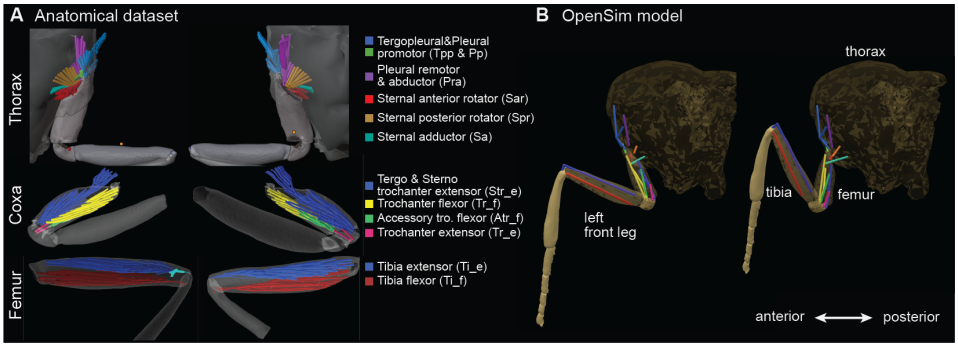

Figure 2: **Front leg muscle reconstructions.** **(A)** Muscles of the thorax, coxa, and femur were segmented from high-resolution X-ray scans (Kuan et al., 2020; Dinges et al., 2021) and visualized within a 3D mesh of the foreleg in Blender. Colors denote anatomically grouped muscles, including biarticular and joint-specific actuators. **(B)** Corresponding muscle-tendon units implemented in OpenSim preserve anatomical attachment points and fiber routing. Color coding is the same as in panel A.

3.2 Muscle model construction and optimization

We modeled 15 muscle-tendon units (MTUs) per foreleg—7 in the thorax, 6 in the coxa, and 2 in the femur—capturing 12 of the 19 muscle groups in our anatomical datasets (Figure 2B). Muscles housed in the tibia were excluded because the tibia segment was partially captured in the X-ray data. In the femur, we modeled the fast tibia flexor and extensor which dominate force generation at the femur–tibia joint (Azevedo et al., 2020). Trochanter muscles were omitted because their function remains unclear (Soler et al., 2004).

Each muscle group was modeled using one or two MTUs, based on a Hill-type formulation (Hill, 1938). This formulation includes a contractile element, a passive parallel elastic component, and a series elastic component, assuming rigid tendons (Millard et al., 2013) (Figure 3A). To keep the model computationally manageable in OpenSim (Delp et al., 2007), we selected up to two representative MTUs per muscle group to capture the main muscle fiber functions (see Supplementary Material for more details). Muscle-tendon lengths and initial attachment points were initialized using anatomical data (Figure 3A, right). We estimated maximum isometric forces from the physiological cross-sectional areas (PCSAs) of the muscles, scaling them with a specific tension of $28\mathrm{mN}/\mathrm{mm}^2$. This value falls between those reported for *Drosophila* jump muscles ($37\mathrm{mN}/\mathrm{mm}^2$) (Eldred et al., 2010) and indirect flight muscles ($9\mathrm{mN}/\mathrm{mm}^2$) (Swank, 2012). We estimated max contraction velocity using a recorded X-ray video of muscle contraction during leg movements.

Because experimental data are limited, initial parameters may not faithfully reflect biological reality. Therefore, we refined parameters through optimization in OpenSim. For each candidate parameter set proposed by the optimizer, static optimization (SO) inferred muscle activations from reference joint angles, and forward dynamics (FD) simulated joint trajectories based on those activations. We used NSGA-II to identify parameter sets that produced kinematics best matching the experimental data (Figure 3B). To avoid overfitting, we optimized parameters across two behaviors—antennal grooming and locomotion—retaining solutions that performed well for both behaviors (Figure 3C). Each joint (1 or 3 DoFs) was optimized independently. For a 3-DoF joint, optimization took approximately 8 h with 200 individuals over 40 generations. The full foreleg optimization (7 DoFs, 15 MTUs) would require $\sim$ 20 hours if each optimization were run sequentially using an Intel i9-14900 processor with 64 GB of RAM.

To assess the biological plausibility of our resulting muscle models, we analyzed muscle moment arms relative to joint motion. Our model correctly reproduced opposing signs for flexor and extensor

moment arms, and predicted the dominant contributors to joint movements (thorax-coxa, and coxa-trochanter yaw, pitch, and roll, and femur-tibia pitch) in agreement with their known functional roles (Figure 3D) (Azevedo et al., 2024; Soler et al., 2004).

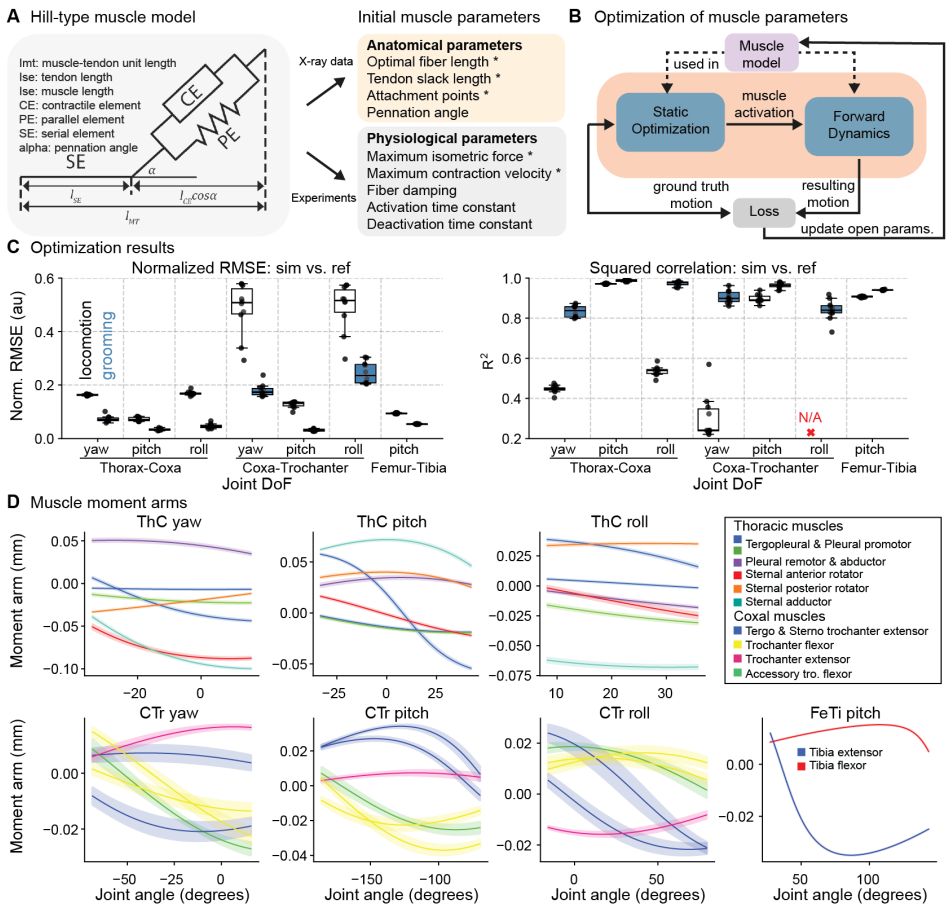

Figure 3: **Optimization and assessment of muscle model parameters.** **(A)** Schematic representation of a Hill-type muscle model. The contractile element (CE) produces active force, the parallel elastic element (PE) provides passive stiffness, and the series elastic element (SE) represents tendon elasticity. l_m is the fiber length, α the pennation angle, and l_t the tendon length. The total muscle-tendon length is $l_{\text{mt}} = l_m \cos(\alpha) + l_t$. Anatomical parameters (fiber length, attachment points, pennation angle) are initialized from X-ray data; physiological parameters (maximum isometric force, activation dynamics) from experiments and literature. Parameters indicated with asterisks are modified through optimization to match measured limb movements. Otherwise parameters were kept at fixed values. **(B)** Optimization pipeline for muscle parameters. Muscle parameters are jointly tuned to minimize the error between measured and simulated joint angles for two behaviors (antennal grooming and locomotion). **(C)** Performance of muscle parameter optimization. Left: root mean squared error (RMSE), normalized with respect to joint movement range. Right: square of the Pearson correlation between simulated and reference joint angles across 7 DoFs (3 joints). Because the coxa-trochanter roll joint does not move during locomotion, it is not shown and has constant values. Data from the top 10 individuals are shown and indicated (dots). **(D)** Moment arms of thoracic, coxal, and femoral muscles with respect to a single joint DoF. Other joints are fixed at mid-range positions. Plots show means and standard deviations across 10 individuals.

3.3 Real limb kinematics data

2D and 3D limb kinematics were measured from tethered flies behaving spontaneously on an air-supported spherical treadmill. Foreleg and head movements were tracked in 2D using DeepLabCut

(Mathis et al., 2018). Separate models were trained for each camera angle. Annotations were refined over multiple iterations to improve accuracy. Five synchronized cameras were calibrated using a ChArUco board, and 3D poses were reconstructed using Anipose (Karashchuk et al., 2021) with Viterbi filtering and spatiotemporal regularization. The resulting 3D pose data were then used to estimate leg and head joint angles via inverse kinematics using SeqIKPy (Özdil et al., 2024b). Behavioral videos were recorded at 100 Hz and interpolated to 500 Hz to ensure simulation stability. We focused on two distinct behaviors, forward walking and antennal grooming, because they involve different patterns of joint coordination in the front legs and can thus test the capacity of our muscle model to reproduce a wider range of movements.

3.4 Imitation learning

To reproduce control strategies that animals use to drive efficient limb movements, we trained policies using imitation learning of reference motion trajectories. Imitation learning has been widely applied in both computer graphics and motor control, providing a powerful framework for learning complex behaviors in simulated agents (Vaxenburg et al., 2025; La Barbera et al., 2021; Chentanez et al., 2018; Peng et al., 2018; Merel et al., 2017; Peng et al., 2020; Hasenclever et al., 2020). Here, we trained neural network models to drive limb movements by activating muscles. Using MyoConverter (Caggiano et al., 2022), we converted our OpenSim model with optimized muscle parameters into a MuJoCo compatible format. We then created the control task in the MuJoCo physics engine (Todorov et al., 2012), using `dm_control` (Tunyasuvunakool et al., 2020), to track the measured kinematics in a muscle-actuated manner.

We trained multilayer perceptron (MLP) policies with Proximal Policy Optimization (PPO) (Schulman et al., 2017) for 15×10^6 steps at a control frequency of 500 Hz, while running the physics engine at 10 kHz to ensure stability. We used PPO due to its well-known stability and robustness in high-dimensional continuous control tasks and based on its prior use in animal imitation learning tasks (Peng et al., 2018; 2020). The MLPs have 2 hidden layers of 512 units followed by a layer of 256 units each and use the ReLU activation function. Policies are trained with a learning rate of 10^{-5} and a discount factor of 0.99, chosen through hyperparameter search. Each training episode was initialized from a random frame in the motion capture sequence and terminated once the end of the clip was reached.

At each timestep, the agent received an observation including joint angles, 3D positions of selected body parts, muscle states, and the remaining time in the clip as described in (La Barbera et al., 2021). The policy produced continuous muscle input levels (i.e., motor neuron activities) within the range of $[0, 1]$ for each muscle.

The reward function was designed to encourage accurate tracking of reference trajectories in Cartesian space, joint angle space, and joint velocity space. At each timestep, the reward was defined as:

$$r_t = \frac{1}{3}\Big[\exp(-w^{\mathrm{p}} d_t^{\mathrm{xpos}}) + \exp(-w^{\mathrm{p}} d_t^{\mathrm{qpos}}) + \exp(-w^{\mathrm{v}} d_t^{\mathrm{qvel}})\Big], \tag{1}$$

where d_t^{xpos}, d_t^{qpos}, and d_t^{qvel} denote the mean Euclidean distances between the simulated body and the reference trajectory in Cartesian position, joint position, and joint velocity, respectively. The reward was clipped to the range $[0, 1]$. The weights were set to $w^{\mathrm{p}} = 5$, and $w^{\mathrm{e}} = 3$. To increase robustness, joint angles were initialized with white noise of variance 0.02. For a 7-DoF leg model with 15 MTUs, imitation learning took approximately 96 hours to complete on an Intel Core i7-12700 processor with 128 GB of RAM.

4 Results

4.1 Muscle synergies during walking and grooming

To investigate how individual muscles contribute to movement, we analyzed activation patterns during forward walking and grooming. Each walking trial contained three stance-swing cycles, and each grooming trial contained three bouts of leg sweeps (Figure 4A). Using OpenSim's static optimization, we estimated joint torques, muscle forces, and activations from our musculoskeletal model for both behaviors.

Among thoracic muscles, the pleural remotor abductor (Pra) was active at stance onset, consistent with its known role in initiating stance motion (Figure 4B, top) (Azevedo et al., 2024). The tergopleural promotor and pleural promotor were elevated during stance-to-swing transitions, driving the coxa forward (Figure 4B, top). During grooming, by contrast, the trochanter flexor and extensor showed rhythmic, overlapping activity, whereas they were out of phase during locomotion (Figure 4B, bottom).

How are muscles acting in fixed groups (i.e., synergies) to produce these movements? To address this question, we applied Non-negative Matrix Factorization (NMF), a widely used method for uncovering low-dimensional structure in muscle activity (Rabbi et al., 2020), to our muscle activities data. Remarkably, just three muscle primitives explained over 90% of the variance (Figure 4C); therefore, we proceeded with these first three primitives. Each primitive exhibited distinct, non-overlapping temporal dynamics (Figure 4D). The first primitive alone captured more than 80%, and aligned with stance onset during locomotion and half of the grooming cycle.

We further examined synergy weights to reveal both invariance and flexibility. Sar and Sa contributed consistently across all synergies and behaviors (Figure 4E), suggesting that they were task-invariant muscles. By contrast, coxal muscles showed task-specific specializations: flexors contributed primarily to synergy 2, whereas extensors contributed to synergy 3 for grooming. These specialization was absent in locomotor primitives (Figure 4E), suggesting that these muscles might have broader roles during locomotion.

Together, these findings predict that muscle coordination is highly behavior-dependent and that the fly can flexibly repurpose the same musculature by engaging distinct, task-specific synergies.

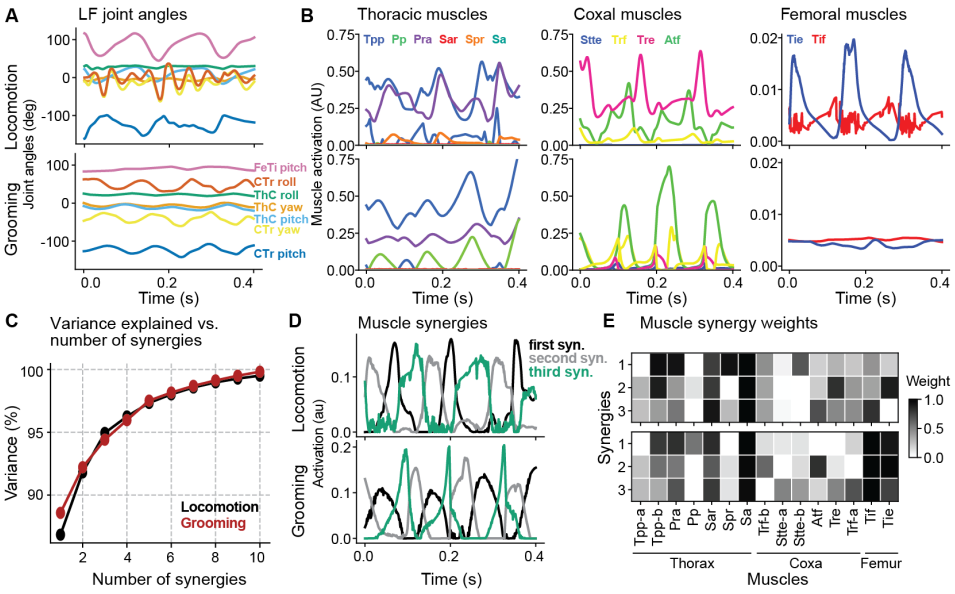

Figure 4: **Predicting muscle synergies from simulated muscle activations in OpenSim. (A)** Simulated joint angle trajectories of the left foreleg (LF) during locomotion (top) and grooming (bottom), obtained from the Static Optimization-Forward Dynamics pipeline. **(B)** Simulated muscle activation dynamics for thoracic (left), coxal (middle), and femoral (right) muscles during locomotion (top) and grooming (bottom). **(C)** Variance explained by increasing numbers of synergies for locomotion (black) and grooming (red). Three primitives capture over 90% of the variance. **(D)** Time courses for the first three extracted synergies (black, gray, and green). **(E)** Synergy weight matrix showing each muscle's contribution; darker values indicate stronger loading.

4.2 Passive joint properties facilitate muscle-driven control

Movement is shaped not only by active muscle forces but also from passive biomechanical forces in the legs that can be either resistive or assistive (Ache and Matheson, 2013). These properties

are thought to offload some control effort from the nervous system, suggesting that they may also facilitate learning in musculoskeletal models.

To test this hypothesis, we systematically varied the passive properties of our MuJoCo model's joints. We considered three parameters: stiffness, which produces spring-like forces that return joints to their reference angles; damping, which resists velocity; and armature, which adds joint inertia (Figure 5A). Armature was included in all conditions to stabilize simulations, while stiffness and damping were selectively added or removed, yielding four different biomechanical models (Figure 5B). Each model was then trained with PPO to imitate locomotor kinematics for 15 million steps. We quantified learning performance by measuring the average reward early (2.5-3M steps) and late (14.5-15M steps) in training (Figure 5C).

We found that the combination of stiffness and damping yields the fastest learning and the highest performance, outperforming models with either property alone (Figure 5D, left). This advantage persisted to the end of training (Figure 5D, right), although final kinematics were qualitatively indistinguishable across models (Figure 5E). We next asked how passive properties shape muscle control. Averaging muscle activations across time and seeds revealed similar mean levels across conditions, but with different temporal dynamics (Figure 5F, Supplementary Material).

Overall, these results support the idea that passive joint mechanics can substantially improve the learnability and robustness of muscle-driven control. We speculate that this is because they stabilize motion and reduce the need for constant correction, allowing the policy to focus on producing effective coordinated muscle patterns rather than correcting instability. In biology, a similar division of labor between passive mechanics and active control might support movement efficiency and robustness (Ache and Matheson, 2013).

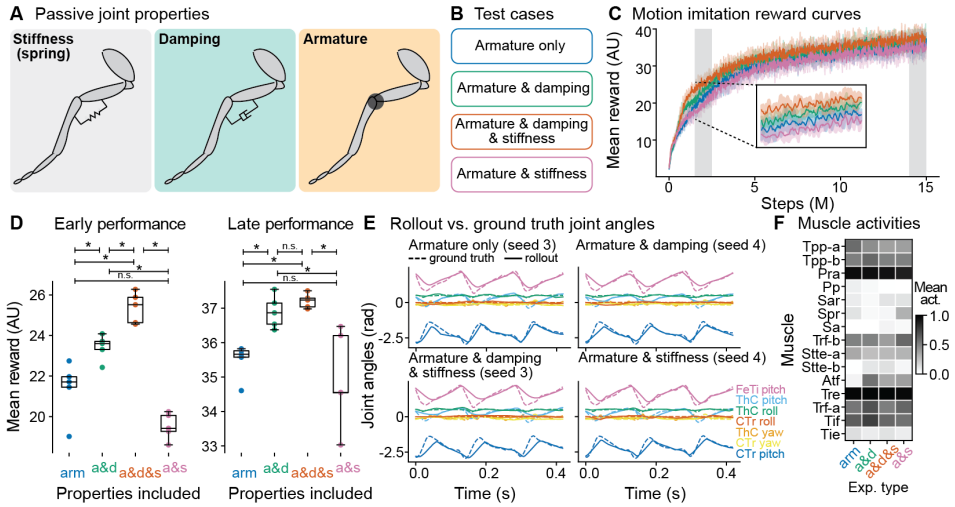

Figure 5: **The impact of passive joint properties on imitation learning of limb kinematics. (A)** Schematic of passive joint properties modified in MuJoCo: stiffness (spring), damping, and armature (inertia). **(B)** Experimental conditions combining these properties: (i) armature only, (ii) armature and damping, (iii) armature and damping and stiffness, and (iv) armature and stiffness. **(C)** Motion imitation rewards at the evaluation stage over the training period, averaged across 5 random seeds (lines) with 5th-95th percentile ranges (shaded regions). Gray boxes indicate evaluation windows for early and late performance in panel D. **(D)** Mean rewards at early (left) and late (right) training stages. Each point is one seed. Statistical comparisons used two-sided Mann-Whitney U tests with Holm-Bonferroni correction (* $P < 0.05$, ns: not significant). **(B-D)** Color code is the same across panels. **(E)** Joint angle trajectories from ground-truth kinematics (dashed) and policy rollouts (solid) for each condition for seeds with the highest final reward. **(F)** Time-averaged muscle activities across conditions; darker shading indicates higher activation.

5 DISCUSSION

Here, we present an end-to-end pipeline for constructing morphologically realistic Hill-type muscle models based on annotated anatomical data. Our framework extracts each muscle's anatomical features, estimates physiological parameters, and optimizes unknown parameters using multi-objective optimization. To our knowledge, a similarly comprehensive and automated approach linking anatomical inputs to functional, muscle-driven simulations is currently lacking in the field, particularly for widely used model organisms like *Drosophila melanogaster*.

We modeled 15 muscle-tendon units (MTUs) per front leg in NeuroMechFly, a biomechanical model of adult *Drosophila* (Lobato-Rios et al., 2022; Wang-Chen et al., 2024). After optimizing muscle parameters, we evaluated how each MTU contributes to different joint DoFs during locomotion and grooming. Our results predict that grooming and locomotion employ distinct muscle synergies. In addition, we show that incorporating passive joint properties —such as damping and stiffness— provides useful priors for musculoskeletal control, improving learning efficiency. Together, these results demonstrate that combining detailed musculature with biologically inspired joint mechanics enables fast, robust, and scalable simulations of muscle-driven behavior and facilitates the efficient training of neural controllers.

Our work provides a critical interface between the outputs of neural network controllers and physical movements. Placing a model of the musculoskeletal system —an additional layer of processing— between the policy network's output and physical actions stabilizes the control task by making the action space better formed and more error-tolerant. The possible actions taken by the policy now sit in a space that, through the musculoskeletal model, can only generate physically plausible movements. This eliminates many unrealistic or even catastrophic actions that would otherwise burden the learning process. In other words, whereas a model that naively controls target joint states would have to first learn (explicitly or implicitly) a world model of physics before learning a policy that operates within physical constraints, our model only needs to learn the latter. This approach enables efficient training at scale, offering a promising path to reduce the sim-to-real gap (Wechsler et al., 2024) without requiring costly task-specific real-world data. Our work focuses on *Drosophila*, an animal with the most complex leg kinematics whose entire nervous system has been mapped. This opens the door to testing neural architectures which are optimized through evolution in embodied settings. We believe our work can impact neuroscience, robotics, and machine learning, areas that share a common focus of understanding and replicating motor control in physically grounded systems.

6 LIMITATIONS AND FUTURE WORK

Our model has several limitations that can be addressed in future work. One major challenge is the scarcity of experimental data and high uncertainty in our chosen physiological parameters. For example, key properties such as the maximum isometric forces and contraction velocities were not directly measured but instead were estimated and optimized using a combination of anatomical and physiological data. Our model will benefit from the acquisition of further experimental data, allowing users to narrow down the space of possible parameter values.

Another limitation is the omission of contact forces from body–body and body–environment interactions. Without modeling external forces, muscle activation patterns may not accurately reflect the demands of untethered behaviors such as locomotion. Incorporating active collision handling will be essential for improving biomechanical realism.

Despite these limitations, our work represents the first biologically-grounded muscle modeling framework for studying motor control in adult *Drosophila melanogaster*. It enables researchers to test hypotheses of muscle function and provides a foundation for uncovering unmodeled passive dynamics and emergent biomechanical properties. Integrating this framework with emerging experimental datasets, such as *in vivo* muscle imaging (Melis et al., 2024), will help refine physiological and anatomical constraints, narrow down the parameter space, and thereby improving the predictive power of musculoskeletal models.

REFERENCES

Donald H Edwards. Neuromechanical simulation. *Frontiers in behavioral neuroscience*, 4:40, 2010.

Jessica Ausborn, Natalia A Shevtsova, and Simon M Danner. Computational modeling of spinal locomotor circuitry in the age of molecular genetics. *International Journal of Molecular Sciences*, 22(13):6835, 2021.

Yongi Kim, Shinya Aoi, Soichiro Fujiki, Simon M Danner, Sergey N Markin, Jessica Ausborn, Ilya A Rybak, Dai Yanagihara, Kei Senda, and Kazuo Tsuchiya. Contribution of afferent feedback to adaptive hindlimb walking in cats: A neuromusculoskeletal modeling study. *Frontiers in Bioengineering and Biotechnology*, 10, 2022.

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.

Pavan Ramdya and Auke Jan Ijspeert. The neuromechanics of animal locomotion: From biology to robotics and back. *Science Robotics*, 8(78):eadg0279, 2023.

Thomas Geijtenbeek, Michiel Van De Panne, and A Frank Van Der Stappen. Flexible muscle-based locomotion for bipedal creatures. *ACM Transactions on Graphics (TOG)*, 32(6):1–11, 2013.

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

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 (TOG)*, 31(4):1–11, 2012.

Diego Aldarondo, Josh Merel, Jesse D Marshall, Leonard Hasenclever, Ugne Klibaite, Amanda Gellis, Yuval Tassa, Greg Wayne, Matthew Botvinick, and Bence P Ölveczky. A virtual rodent predicts the structure of neural activity across behaviors. *Nature*, pages 1–3, 2024.

Sibo Wang-Chen, Victor Alfred Stimpfling, Thomas Ka Chung Lam, Pembe Gizem Özdil, Louise Genoud, Femke Hurtak, and Pavan Ramdya. Neuromechfly v2: simulating embodied sensorimotor control in adult drosophila. *Nature Methods*, pages 1–10, 2024.

Hugo J Bellen, Chao Tong, and Hiroshi Tsuda. 100 years of drosophila research and its impact on vertebrate neuroscience: a history lesson for the future. *Nature Reviews Neuroscience*, 11(7): 514–522, 2010.

Gerald M Rubin and Allan C Spradling. Genetic transformation of drosophila with transposable element vectors. *Science*, 218(4570):348–353, 1982.

Sven Dorkenwald, Arie Matsliah, Amy R Sterling, Philipp Schlegel, Szi-Chieh Yu, Claire E McKellar, Albert Lin, Marta Costa, Katharina Eichler, Yijie Yin, et al. Neuronal wiring diagram of an adult brain. *Nature*, 634(8032):124–138, 2024.

Jasper S. Phelps, David Grant Colburn Hildebrand, Brett J. Graham, Aaron T. Kuan, Logan A. Thomas, Tri M. Nguyen, Julia Buhmann, Anthony W. Azevedo, Anne Sustar, Sweta Agrawal, Mingguan Liu, Brendan L. Shanny, Jan Funke, John C. Tuthill, and Wei-Chung Allen Lee. Reconstruction of motor control circuits in adult Drosophila using automated transmission electron microscopy. *Cell*, 184(3):759–774.e18, 2021. ISSN 0092-8674. doi: 10.1016/j.cell.2020.12.013.

Sophie Aimon, Takeo Katsuki, Tongqiu Jia, Logan Grosenick, Michael Broxton, Karl Deisseroth, Terrence J Sejnowski, and Ralph J Greenspan. Fast near-whole–brain imaging in adult drosophila during responses to stimuli and behavior. *PLoS biology*, 17(2):e2006732, 2019.

Julie H Simpson. Descending control of motor sequences in drosophila. *Current opinion in neurobiology*, 84:102822, 2024.

Chin-Lin Chen, Laura Hermans, Meera C Viswanathan, Denis Fortun, Florian Aymanns, Michael Unser, Anthony Cammarato, Michael H Dickinson, and Pavan Ramdya. Imaging neural activity in the ventral nerve cord of behaving adult Drosophila. *Nature communications*, 9(1), 2018.

Márcia M Aranha and Maria Luísa Vasconcelos. Deciphering drosophila female innate behaviors. *Current opinion in neurobiology*, 52:139–148, 2018.

Pembe Gizem Özdil, Jonathan Arreguit, Clara Scherrer, Auke Ijspeert, and Pavan Ramdya. Centralized brain networks underlie body part coordination during grooming. *bioRxiv*, 2024a. doi: 10.1101/2024.12.17.628844.

Roman Vaxenburg, Igor Siwanowicz, Josh Merel, Alice A Robie, Carmen Morrow, Guido Novati, Zinovia Stefanidi, Gert-Jan Both, Gwyneth M Card, Michael B Reiser, et al. Whole-body physics simulation of fruit fly locomotion. *Nature*, pages 1–3, 2025.

Victor Lobato-Rios, Shravan Tata Ramalingasetty, Pembe Gizem Özdil, Jonathan Arreguit, Auke Jan Ijspeert, and Pavan Ramdya. NeuroMechFly, a neuromechanical model of adult Drosophila melanogaster. *Nature Methods*, 19(5):620–627, May 2022. ISSN 1548-7105. doi: 10.1038/s41592-022-01466-7.

Anthony Azevedo, Ellen Lesser, Jasper S Phelps, Brandon Mark, Leila Elabbady, Sumiya Kuroda, Anne Sustar, Anthony Moussa, Avinash Khandelwal, Chris J Dallmann, et al. Connectomic reconstruction of a female drosophila ventral nerve cord. *Nature*, pages 1–9, 2024.

Ellen Lesser, Anthony W Azevedo, Jasper S Phelps, Leila Elabbady, Andrew Cook, Durafshan Sakeena Syed, Brandon Mark, Sumiya Kuroda, Anne Sustar, Anthony Moussa, et al. Synaptic architecture of leg and wing premotor control networks in drosophila. *Nature*, 631(8020): 369–377, 2024.

Gesa F. Dinges, Alexander S. Chockley, Till Bockemühl, Kei Ito, Alexander Blanke, and Ansgar Büschges. Location and arrangement of campaniform sensilla in Drosophila melanogaster. *Journal of Comparative Neurology*, 529(4):905–925, 2021. ISSN 1096-9861. doi: 10.1002/cne.24987.

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

Jeongseok Lee, Michael X. Grey, Sehoon Ha, Tobias Kunz, Sumit Jain, Yuting Ye, Siddhartha S. Srinivasa, Mike Stilman, and C Karen Liu. Dart: Dynamic animation and robotics toolkit. *The Journal of Open Source Software*, 3(22):500, 2018b.

Thomas Geijtenbeek. The hyfydy simulation software. 2021.

Seungmoon Song, Łukasz Kidziński, Xue Bin Peng, Carmichael Ong, Jennifer Hicks, Sergey Levine, Christopher G Atkeson, and Scott L Delp. Deep reinforcement learning for modeling human locomotion control in neuromechanical simulation. *Journal of neuroengineering and rehabilitation*, 18:1–17, 2021.

Vittorio Caggiano, Guillaume Durandau, Huiyi Wang, Chun Kwang Tan, Pierre Schumacher, Huawei Wang, Alberto Silvio Chiappa, Alessandro Marin Vargas, Alexander Mathis, Jungdam Won, et al. Myochallenge 2023: Towards human-level dexterity and agility. 2024.

Adam D Sylvester, Steven G Lautzenheiser, and Patricia Ann Kramer. A review of musculoskeletal modelling of human locomotion. *Interface Focus*, 11(5):20200060, 2021.

Ajay Seth, Jennifer L Hicks, Thomas K Uchida, Ayman Habib, Christopher L Dembia, James J Dunne, Carmichael F Ong, Matthew S DeMers, Apoorva Rajagopal, Matthew Millard, et al. OpenSim: Simulating musculoskeletal dynamics and neuromuscular control to study human and animal movement. *PloS computational biology*, 14(7), 2018.

James P Charles, Ornella Cappellari, and John R Hutchinson. A dynamic simulation of musculoskeletal function in the mouse hindlimb during trotting locomotion. *Frontiers in Bioengineering and Biotechnology*, 6:61, 2018.

Shravan Tata Ramalingasetty, Simon M. Danner, Jonathan Arreguit, Sergey N. Markin, Dimitri Rodarie, Claudia Kathe, Gregoire Courtine, Ilya A. Rybak, and Auke Jan Ijspeert. A Whole-Body Musculoskeletal Model of the Mouse. *IEEE Access*, 9:163861–163881, 2021. ISSN 2169-3536. doi: 10.1109/ACCESS.2021.3133078. URL https://ieeexplore.ieee.org/document/9638502/.

Jesse I Gilmer, Susan K Coltman, Geraldine Cuenu, John R Hutchinson, Daniel Huber, Abigail L Person, and Mazen Al Borno. A novel biomechanical model of the proximal mouse forelimb predicts muscle activity in optimal control simulations of reaching movements. *Journal of Neurophysiology*, 133(4):1266–1278, 2025.

Travis DeWolf, Steffen Schneider, Paul Soubiran, Adrian Roggenbach, and Mackenzie Weygandt Mathis. Neuro-musculoskeletal modeling reveals muscle-level neural dynamics of adaptive learning in sensorimotor cortex. *bioRxiv*, pages 2024–09, 2024.

Naomichi Ogihara, Haruyuki Makishima, Shinya Aoi, Yasuhiro Sugimoto, Kazuo Tsuchiya, and Masato Nakatsukasa. Development of an anatomically based whole-body musculoskeletal model of the japanese macaque (macaca fuscata). *American Journal of Physical Anthropology: The Official Publication of the American Association of Physical Anthropologists*, 139(3):323–338, 2009.

Sherwin S Chan and Daniel W Moran. Computational model of a primate arm: from hand position to joint angles, joint torques and muscle forces. *Journal of neural engineering*, 3(4):327, 2006.

Pasha A van Bijlert, Thomas Geijtenbeek, Ineke H Smit, Anne S Schulp, and Karl T Bates. Muscle-driven predictive physics simulations of quadrupedal locomotion in the horse. *Integrative and Comparative Biology*, 64(3):694–714, 2024.

Jeffery W Rankin, Jonas Rubenson, and John R Hutchinson. Inferring muscle functional roles of the ostrich pelvic limb during walking and running using computer optimization. *Journal of the Royal Society Interface*, 13(118):20160035, 2016.

Vittorio La Barbera, Fabio Pardo, Yuval Tassa, Monica Daley, Christopher Richards, Petar Kormushev, and John Hutchinson. Ostrichrl: A musculoskeletal ostrich simulation to study biomechanical locomotion. *arXiv preprint arXiv:2112.06061*, 2021.

Mengdi Zhao, Ning Wang, Xinrui Jiang, Xiaoyang Ma, Haixin Ma, Gan He, Kai Du, Lei Ma, and Tiejun Huang. An integrative data-driven model simulating c. elegans brain, body and environment interactions. *Nature Computational Science*, 4(12):978–990, 2024.

Shihui Guo, Juncong Lin, Toni Wöhrl, and Minghong Liao. A Neuro-Musculo-Skeletal Model for Insects With Data-driven Optimization. *Scientific reports*, 8(1):1–11, 2018.

Jonathan Arreguit, Shravan Tata Ramalingasetty, and Auke Ijspeert. Farms: Framework for animal and robot modeling and simulation. *bioRxiv*, 2023.

Emilio Bizzi and Vincent CK Cheung. The neural origin of muscle synergies. *Frontiers in computational neuroscience*, 7:51, 2013.

Lena H Ting and J Lucas McKay. Neuromechanics of muscle synergies for posture and movement. *Current opinion in neurobiology*, 17(6):622–628, 2007.

Akira Saito, Aya Tomita, Ryosuke Ando, Kohei Watanabe, and Hiroshi Akima. Muscle synergies are consistent across level and uphill treadmill running. *Scientific reports*, 8(1):5979, 2018.

Matthew C Tresch, Vincent CK Cheung, and Andrea d'Avella. Matrix factorization algorithms for the identification of muscle synergies: evaluation on simulated and experimental data sets. *Journal of neurophysiology*, 95(4):2199–2212, 2006.

Simon Sponberg, Thomas L Daniel, and Adrienne L Fairhall. Dual dimensionality reduction reveals independent encoding of motor features in a muscle synergy for insect flight control. *PLoS computational biology*, 11(4):e1004168, 2015.

Mazen Al Borno, Jennifer L Hicks, and Scott L Delp. The effects of motor modularity on performance, learning and generalizability in upper-extremity reaching: a computational analysis. *Journal of the Royal Society Interface*, 17(167):20200011, 2020.

Cameron Berg, Vittorio Caggiano, and Vikash Kumar. Sar: Generalization of physiological agility and dexterity via synergistic action representation. *Autonomous Robots*, 48(8):28, 2024.

Salil S. Bidaye, Till Bockemühl, and Ansgar Büschges. Six-legged walking in insects: how CPGs, peripheral feedback, and descending signals generate coordinated and adaptive motor rhythms. *Journal of Neurophysiology*, 119(2):459–475, 2018. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.00658.2017.

Sasha N Zill, Chris J Dallmann, Ansgar Bueschges, Sumaiya Chaudhry, and Josef Schmitz. Force dynamics and synergist muscle activation in stick insects: the effects of using joint torques as mechanical stimuli. *Journal of neurophysiology*, 120(4):1807–1823, 2018.

Durafshan Sakeena Syed, Primoz Ravbar, and Julie H Simpson. Inhibitory circuits coordinate leg movements during drosophila grooming. *bioRxiv*, pages 2024–06, 2024.

Moritz Haustein, Alexander Blanke, Till Bockemühl, and Ansgar Büschges. A leg model based on anatomical landmarks to study 3d joint kinematics of walking in drosophila melanogaster. *Frontiers in Bioengineering and Biotechnology*, 12:1357598, 2024.

Cédric Soler, Malgorzata Daczewska, Jean Philippe Da Ponte, Bernard Dastugue, and Krzysztof Jagla. Coordinated development of muscles and tendons of the Drosophila leg. *Development*, 131 (24):6041–6051, 2004.

Douglas M Swank. Mechanical analysis of drosophila indirect flight and jump muscles. *Methods*, 56(1):69–77, 2012.

Jan M Ache and Thomas Matheson. Passive joint forces are tuned to limb use in insects and drive movements without motor activity. *Current Biology*, 23(15):1418–1426, 2013.

Aaron T Kuan, Jasper S Phelps, Logan A Thomas, Tri M Nguyen, Julie Han, Chiao-Lin Chen, Anthony W Azevedo, John C Tuthill, Jan Funke, Peter Cloetens, et al. Dense neuronal reconstruction through X-ray holographic nano-tomography. *Nature neuroscience*, 2020.

Anthony W Azevedo, Evyn S Dickinson, Pralaksha Gurung, Lalanti Venkatasubramanian, Richard S Mann, and John C Tuthill. A size principle for recruitment of Drosophila leg motor neurons. *eLife*, 9:e56754, 2020.

Archibald Vivian 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.

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.

Catherine C Eldred, Dimitre R Simeonov, Ryan A Koppes, Chaoxing Yang, David T Corr, and Douglas M Swank. The mechanical properties of drosophila jump muscle expressing wild-type and embryonic myosin isoforms. *Biophysical journal*, 98(7):1218–1226, 2010.

Alexander Mathis, Pranav Mamidanna, Kevin M. Cury, Taiga Abe, Venkatesh N. Murthy, Mackenzie Weygandt Mathis, and Matthias Bethge. DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nature Neuroscience*, 21(9):1281–1289, 2018. ISSN 1097-6256, 1546-1726. doi: 10.1038/s41593-018-0209-y.

Pierre Karashchuk, Katie L. Rupp, Evyn S. Dickinson, Sarah Walling-Bell, Elischa Sanders, Eiman Azim, Bingni W. Brunton, and John C. Tuthill. Anipose: A toolkit for robust markerless 3D pose estimation. *Cell Reports*, 36(13):109730, 2021. ISSN 22111247. doi: 10.1016/j.celrep.2021.109730.

Pembe Gizem Özdil, Auke Ijspeert, and Pavan Ramdya. SeqIKPy: a python package for inverse kinematics in insects, 2024b. URL `https://doi.org/10.5281/zenodo.12601317`.

Nuttapong Chentanez, Matthias Müller, Miles Macklin, Viktor Makoviychuk, and Stefan Jeschke. Physics-based motion capture imitation with deep reinforcement learning. In *Proceedings of the 11th ACM SIGGRAPH Conference on Motion, Interaction and Games*, pages 1–10, 2018.

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):1–14, July 2018. ISSN 1557-7368. doi: 10.1145/3197517.3201311. URL http://dx.doi.org/10.1145/3197517.3201311.

Josh Merel, Yuval Tassa, Dhruva TB, Sriram Srinivasan, Jay Lemmon, Ziyu Wang, Greg Wayne, and Nicolas Heess. Learning human behaviors from motion capture by adversarial imitation. *arXiv preprint arXiv:1707.02201*, 2017.

Xue Bin Peng, Erwin Coumans, Tingnan Zhang, Tsang-Wei Lee, Jie Tan, and Sergey Levine. Learning agile robotic locomotion skills by imitating animals. *arXiv preprint arXiv:2004.00784*, 2020.

Leonard Hasenclever, Fabio Pardo, Raia Hadsell, Nicolas Heess, and Josh Merel. CoMic: Complementary task learning amp; mimicry for reusable skills. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 4105–4115. PMLR, 13–18 Jul 2020. URL https://proceedings.mlr.press/v119/hasenclever20a.html.

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, 2012. doi: 10.1109/IROS.2012.6386109. ISSN: 2153-0866.

Saran Tunyasuvunakool, Alistair Muldal, Yotam Doron, Siqi Liu, Steven Bohez, Josh Merel, Tom Erez, Timothy Lillicrap, Nicolas Heess, and Yuval Tassa. dm_control: Software and tasks for continuous control. *Software Impacts*, 6:100022, 2020.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017.

Mohammad Fazle Rabbi, Claudio Pizzolato, David G Lloyd, Chris P Carty, Daniel Devaprakash, and Laura E Diamond. Non-negative matrix factorisation is the most appropriate method for extraction of muscle synergies in walking and running. *Scientific reports*, 10(1):8266, 2020.

Iris Wechsler, Alexander Wolf, Julian Shanbhag, Sigrid Leyendecker, Bjoern M. Eskofier, Anne D. Koelewijn, Sandro Wartzack, and Jörg Miehling. Bridging the sim2real gap. investigating deviations between experimental motion measurements and musculoskeletal simulation results—a systematic review. *Frontiers in Bioengineering and Biotechnology*, 12, June 2024. ISSN 2296-4185. doi: 10.3389/fbioe.2024.1386874. URL http://dx.doi.org/10.3389/fbioe.2024.1386874.

Johan M Melis, Igor Siwanowicz, and Michael H Dickinson. Machine learning reveals the control mechanics of an insect wing hinge. *Nature*, pages 1–9, 2024.

A Supplementary Material

A.1 Muscle reconstruction and modeling of the middle and hind legs

We concentrated our efforts on developing a fully functional front-leg muscle model for two main reasons: (i) we have access to a rich kinematic dataset spanning diverse behaviors such as grooming and locomotion, and (ii) multiple complementary datasets exist that validate joint rotation centers and degrees of freedom (DoFs).

Nevertheless, we also annotated muscle fibers in the midlegs (Figure S1A) and hindlegs (Figure S2A) using our custom X-ray dataset. In total, we modeled seven muscle–tendon units (MTUs) per midleg (Figure S1B) and eight MTUs per hindleg (Figure S2B). Because only a single dataset was available for these legs, our ability to examine muscle attachment points across joint configurations or to cross-validate the reconstructions was limited. For this reason, we restricted our work to anatomical reconstruction rather than optimizing the models to reproduce motor behaviors.

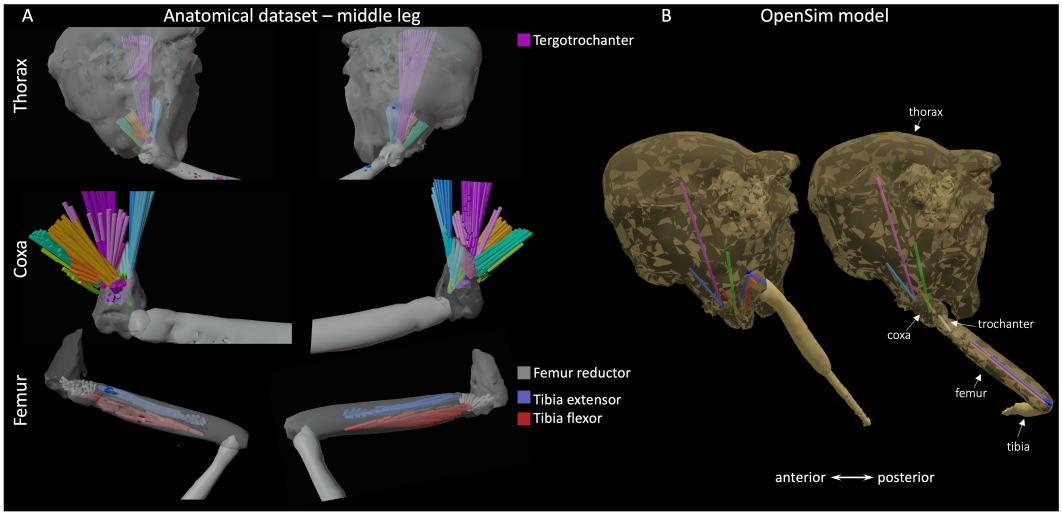

Figure S1: **Middle leg muscle reconstructions**. **(A)** Muscles of the thorax, coxa, and femur were segmented from a high-resolution X-ray scan and visualized within a 3D mesh of the leg in Blender. Distinct colors denote anatomically grouped muscles.**(B)** Corresponding muscle-tendon units implemented in OpenSim, preserve anatomical attachment points and fiber routing. Color coding is the same as in panel A and reflects functionally grouped muscles.

Future datasets capturing the mid- and hindlegs at high resolution will enable more detailed biomechanical modeling and muscle reconstructions. Once such data become available, our framework can be readily extended to achieve muscle-driven control of all six legs.

A.2 Refining the biomechanical model

A fundamental challenge in muscle modeling is accurately grounding muscle attachment points and joint movements in anatomical data. To better align our model with biological reality, we replaced NeuroMechFly's foreleg meshes with those derived from X-ray scans (Figure S3A), where muscles were reconstructed. This adjustment was motivated by two key factors. First, using the same leg mesh as the dataset ensured a one-to-one mapping between the muscles in the dataset and our model. Second, we identified a major discrepancy in the fully flexed resting posture of the original model's leg compared to real fly anatomy. Specifically, in the experimental data, the trochanter segment of the front leg forms a bridge between the coxa and femur, positioning them side by side rather than stacked vertically when fully flexed (Figure S3B). As a result, when fully extended, the foreleg adopts a parenthesis-like shape rather than forming a straight line (Figure S3A).

With this improved biomechanical model, we next investigated the joint rotational axes and centers that could more accurately replicate real leg movements. To determine these parameters, we exam-

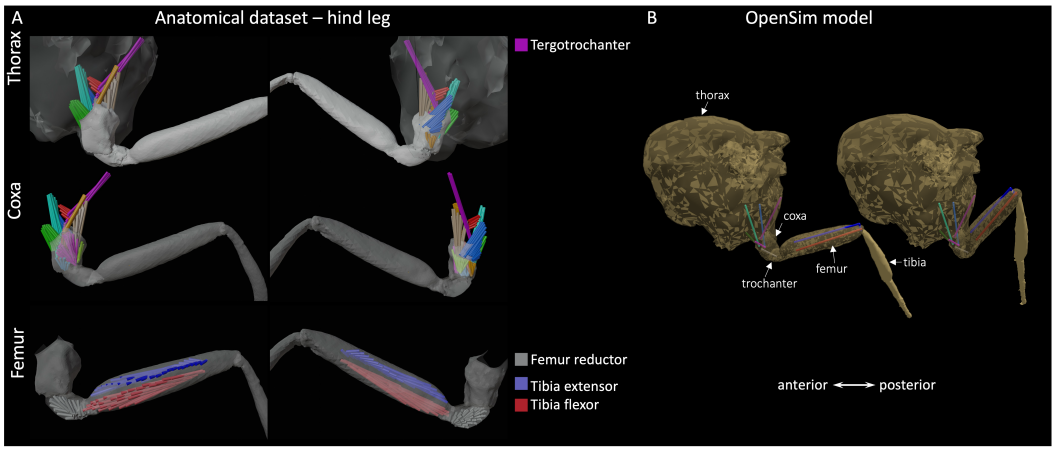

Figure S2: **Hind leg muscle reconstructions.** **(A)** Muscles of the thorax, coxa, and femur were segmented from a high-resolution X-ray scan and visualized within a 3D mesh of the leg in Blender. Distinct colors denote anatomically grouped muscles. **(B)** Corresponding muscle-tendon units implemented in OpenSim, preserve anatomical attachment points and fiber routing. Color coding is the same as in panel A and reflects functionally grouped muscles.

ined muscle attachment points and joint condyles (Haustein et al., 2024). Our analysis suggested that the coxa-trochanter joint may, in fact, possess three DoFs, in contrast to the originally assumed two DoFs (Lobato-Rios et al., 2022).

To test this hypothesis, we used SeqIKPy (Özdil et al., 2024) to design a kinematic chain with three DoFs at the coxa-trochanter joint. The revised joint configuration was equally effective at tracking the leg trajectory compared to the original two-DoF model (Figure S3C). However, notably, the three-DoF configuration reduced the reliance on thorax-coxa joint rotations during antennal grooming, decreasing the range of motion from $[-95°, 50°]$ to a more constrained interval of $[-30°, -7°]$ (Figure S3D), preventing unnatural rotations of the thorax-coxa joint.

A.3 Hill-type muscle model

We chose to use a Hill-type muscle model because it offers a good balance between biological realism and computational efficiency. Hill-type muscle models are widely adopted in the biomechanics and neuromechanics communities due to their versatility and biological plausibility (Caillet et al., 2022). By contrast, Ekeberg-type muscle models are computationally very efficient yet have limited biological correspondence (Ekeberg et al., 2004). Compared with simpler torque-driven or mass-spring systems, Hill-type models explicitly represent the key components of a muscle-tendon unit (MTU)—namely, contractile, series elastic, and parallel elastic elements. This structure allows us to capture essential nonlinear properties of biological muscles, including force-velocity and force-length relationships, which are critical for simulating realistic motor behavior.

We adopted a Hill-type model based on (Millard et al., 2013) and (Geyer et al., 2003), consisting of a Contractile Element (CE) for active force generation, Parallel Elasticity (PE) for passive stiffness, and Series Elasticity (SE) representing the tendons. The Buffer Elasticity (BE) was omitted due to its negligible contribution to muscle dynamics (Geyer and Herr, 2010).

The total muscle force is given by:

$$F_{MT} = (F_{CE} + F_{PE} + F_{\text{damper}}) \cos \alpha = F_{SE} \tag{S1}$$

where F_{CE}, F_{PE}, and F_{damper} are the active, passive, and damping forces, respectively. The active force is computed as:

$$F_{CE} = a(t) F_{\max} f_l(l_{CE}) f_v(v_{CE}) \tag{S2}$$

where $a(t)$ is the activation state, $F_{\max}$ the maximum isometric force, and $f_l(l_{CE})$ and $f_v(v_{CE})$ the force-length and force-velocity relationships, respectively. The passive force F_{PE} is only active during muscle elongation, while the damping force depends on contractile element velocity.

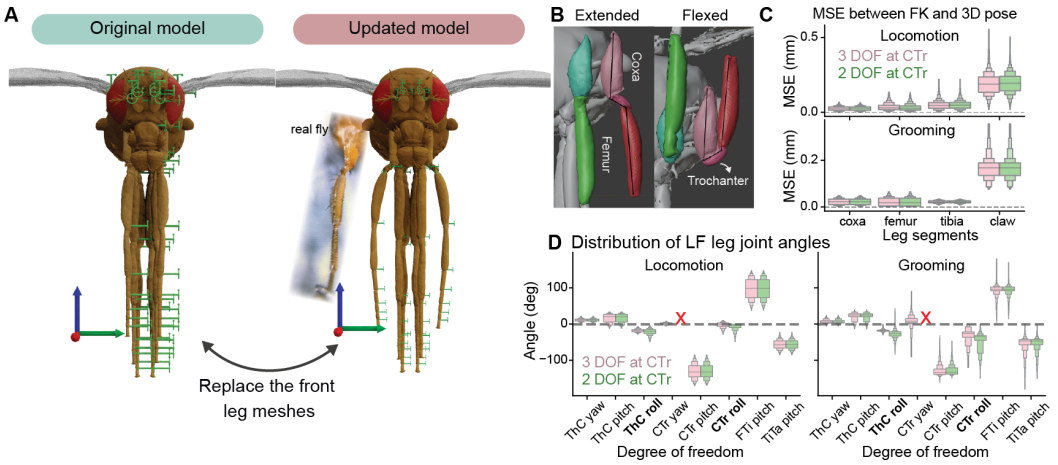

Figure S3: **Enhancements in the biomechanical model of NeuroMechFly.** **(A)** The forelegs of NeuroMechFly were updated by replacing the original meshes with an anatomically detailed muscle-based dataset. Unlike the previous model, where the fully extended leg was nearly straight, the updated model exhibits a natural parenthesis-like curvature, consistent with the biological data (middle). **(B)** The trochanter's role in leg posture. The trochanter bridges the coxa and femur, positioning them side by side rather than in a linear alignment. In the previous model (green), this structural feature was simplified, whereas the updated model (red) better captures the fly's natural joint configuration. **(C)** Mean squared error (MSE) between measured 3D poses and forward kinematics (FK) reconstructed from inverse kinematics, comparing 2-DoF and 3-DoF coxa-trochanter joint configurations during locomotion and grooming. Both kinematic configurations perform equally well. **(D)** Distribution of left foreleg joint angles during locomotion and grooming for different degrees of freedom, showing the effect of anatomical refinements and DoF addition on the joint angle range. ThC roll and CTr roll (in bold) show a visible decrement in range in the 3 DoF model.

For compatibility across simulation engines, we set the pennation angle to zero and scaled $F_{\max}$ accordingly. Assuming a rigid tendon, the model state variables reduce to the muscle fiber length l_{CE}, ensuring C^2-continuity. The complete model is defined by physiological parameters (e.g., $F_{\max}, v_{\max}, \alpha, \tau_{\text{act}}, \tau_{\text{deact}}$) and anatomical parameters (e.g., $l_{\text{opt}}, l_{\text{slack}}$, muscle attachment points, and muscle paths), detailed in Table 1.

We modeled each functional muscle group as one MTU. Unlike human muscles, many fly muscles attach directly to the inner cuticle, often without tendons, and fibers within a group can vary considerably. To model each group, we selected representative fibers using the following criteria: (i) If the muscle attaches via a tendon, we used the tendon's location. If attachment is via multiple fibers, we selected the fiber closest to the group's center of mass—typically also the median-length fiber. (ii) For groups with widely spread attachment points, we subdivided them into smaller, more homogeneous clusters and selected one representative fiber per cluster.

A.4 Muscle modeling in OpenSim

Muscles were defined as *forceset* objects in OpenSim, serving as force generators. The Millard 2013 muscle model (Millard et al., 2013) was used, with parameters categorized into three types (Table 2): (i) open parameters requiring optimization, (ii) fixed parameters set by model choice, and (iii) default parameters left unchanged.

Most anatomical and physiological parameters were open parameters, extracted or estimated through data-driven methods described in the methods section of this paper. Assuming rigid tendons, tendon compliance was ignored, and pennation angles were set to zero, meaning MTU length changes were solely due to muscle fiber length changes.

Control and activation parameters were adjusted for greater flexibility, with a minimum activation of 0.01%, while max control and detailed F-L and F-V curve parameters were left at their default

Table 1: A collection of model parameters used in the muscle model. Letters A and P stand for anatomical and physiological parameters, respectively.

Parameter	Type	Description
Optimal fiber length	A	The fiber length at which actin–myosin interactions are maximal, i.e., when maximum force can be produced.
Tendon slack length	A	The length beyond which a muscle's tendons begin resisting stretch and producing force.
Pennation angles	A	The angle between a fascicle's orientation and the tendon axis.
Muscle cross-section area	A	The cross-section area of muscle fibers.
Muscle attachment points	A	The points to which a muscle tendon unit connects at its fiber start (origin) and tendon end (insertion). Optionally an additional attachment point can be defined for more complex muscle paths
Muscle paths	A	The path that a muscle tendon unit travels along from its origin to its insertion.
Maximum isometric force	P	The maximum isometric force generated by the Contractile Element (CE) at its optimal length l_{opt}.
Maximum contraction velocity	P	The maximum velocity at which the Contractile Element (CE) can contract.
Activation time constants	P	A time constant determining the transfer speed from neural signal to muscle activation.
Deactivation time constants	P	A time constant determining the transfer speed from neural signal to muscle deactivation.
Muscle damping factors	P	The coefficient for the muscle damping force as defined in Millard's muscle equation.

values. If not explicitly defined, these curves followed empirical formulations (Millard et al., 2013). Due to the lack of measured *Drosophila* muscle curves, default curves in OpenSim were assumed to approximate real physiological behavior.

A.5 Muscle parameter optimization using NSGA-II in OpenSim

To optimize muscle parameters, we used NSGA-II, an elitist genetic algorithm well-suited for multi-objective optimization problems (Deb et al., 2002), implemented in Python using the `geatpy` package.

Each muscle–tendon unit (MTU) was optimized in a 6-dimensional parameter space (or 9-dimensional when an additional via point was included in the muscle path), with separate objective metrics defined for each active degree of freedom (DoF). These objectives aimed to minimize both the mean squared error and the reverse correlation between the ground-truth and simulated motion. To reduce overfitting of muscle parameters, we ran the SO-FD pipeline for both locomotion and grooming behaviors (Figure S4). The total objective score was computed by summing the two scores for each DoF. NSGA-II hyperparameters were empirically tuned (Table 3), and parameter search ranges were iteratively refined based on optimization performance. The distribution of the resulting muscle parameter values is shown in Figure S5.

Specifically, the following parameters were optimized:

- Maximum isometric force was computed as the product of a fixed base tension (from prior experimental work), a scaling factor (optimized between 0.3–3), and the physiological cross-sectional area (PCSA) calculated from CT scans.
- Maximum contraction velocity was defined as a base value (estimated from femur–tibia motion videos under X-ray) scaled by a factor between 0.4 and 2.4.

Table 2: Millard 2013 muscle model parameters categorized by type and optimization.

Parameter	Value Type	Parameter Type
Optimal fiber length	Anatomical	Open
Tendon slack length	Anatomical	Open
Geometry path	Anatomical	Open
Max isometric force	Physiological	Open
Max contraction velocity	Physiological	Open
Activation time constant	Physiological	Fixed
Deactivation time constant	Physiological	Fixed
Ignore tendon compliance	Physiological	Fixed
Pennation angle at optimal	Anatomical	Fixed
Maximum pennation angle	Anatomical	Fixed
Min activation	Physiological	Fixed
Min control	Physiological	Default
Max control	Physiological	Default
Fiber damping	Physiological	Default
Active force–length curve	Physiological	Default
Force–velocity curve	Physiological	Default
Fiber force–length curve	Physiological	Default
Tendon force–length curve	Physiological	Default

- Optimal fiber length and tendon slack length were based on ratios observed in the CT data, scaled by a factor between 0.8 and 1.2, and capped at a maximum of 95% of the total length.
- Muscle paths were allowed to vary within a cube of 5-10μm around the annotated insertion points.

Table 3: NSGA-II parameters used in optimization.

Muscles	Generation	Population number	Mutation	Cross-over
Thorax	200	120	0.7	0.5
Coxa	200	40	0.7	0.5
Femur	200	300	0.5	0.3

We applied a curriculum learning strategy to each joint using front leg kinematics during grooming and locomotion. The warm-up stage (5–10 generations) used conservative hyperparameters to encourage rapid convergence, followed by a fine-tuning stage (5–10 generations) that employed more exploratory search settings. Final optimized muscle parameters were validated in the SO-FD pipeline, and the best-performing set was mirrored to the right leg, completing the musculoskeletal model of the front legs.

A.6 Imitation learning using PPO in MuJoCo

We trained multilayer perceptron (MLP) policies with Proximal Policy Optimization (PPO) (Schulman et al., 2017) for 15×10^6 steps at a control frequency of 500 Hz, while running the physics engine at 10,000 Hz to ensure stability. The hyperparameters used for training are listed in Table 4.

Initially, our motion capture dataset included only joint angles and velocities. To enrich this dataset, we replayed the joint kinematics using a point-torque model and recorded all observable physical quantities, including body positions and velocities in Cartesian space. This expanded dataset was then used as the reference trajectory in our imitation learning experiments.

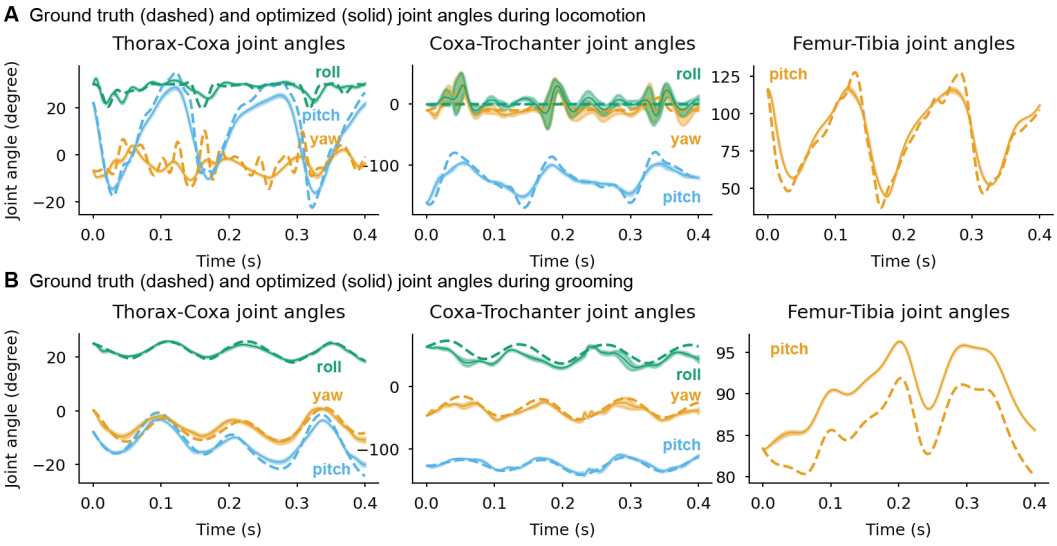

Figure S4: **Joint kinematics across all degrees of freedom in reference and simulated data.** Joint angle trajectories from **(A)** locomotion and **(B)** grooming behaviors are used as reference inputs for optimizing unknown muscle parameters. The MSE and correlation between the reference and resulting trajectories are used as objectives. The thick solid line represents the reference trajectory, the thin solid line shows the trajectory of the best-performing individual from the optimization, and the shaded region indicates the standard deviation across the top 10 individuals.

At each timestep, the agent received an observation vector comprising joint angles, 3D positions of selected body parts (i.e., tibia-tarsus and claw), muscle states (i.e., length, velocity, activation, and force), and the remaining time in the clip, as described in (La Barbera et al., 2021). The policy output consisted of continuous muscle input values in the range $[0, 1]$ for each muscle.

The reward function was designed to encourage accurate tracking of the reference motion, both in joint space and Cartesian space. We initially tested the reward function and training setup using in a point-torque model. Upon achieving successful performance in this simpler setting, we transitioned to the more complex musculoskeletal model. All parameter values were selected through a hyperparameter search.

Table 4: Key hyperparameters used for imitation learning with PPO. The same values were applied to both the actor and critic networks.

Hyperparameter	**Value**
Network Architecture	[512, 512, 256]
Activation Function	ReLU
Optimizer	Adam
Learning Rate	1×10^{-5}
Batch Size	64
Rollout Length (n_{steps})	2048
Epochs per Update	10

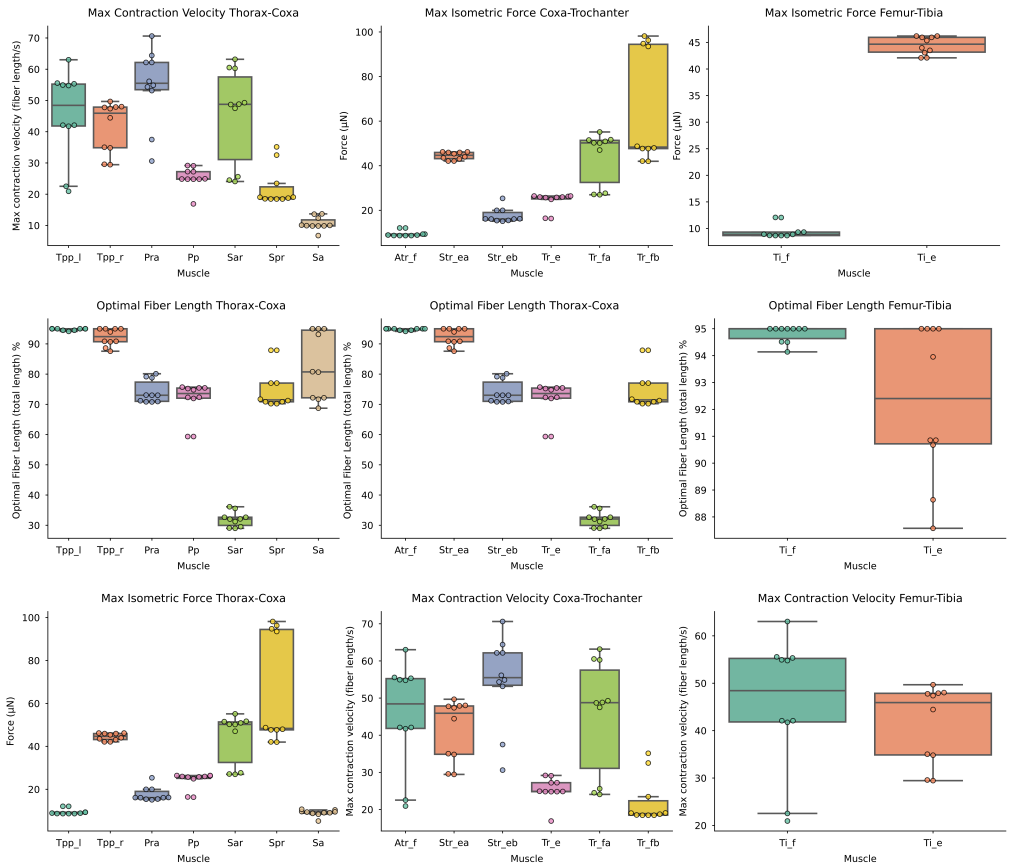

Figure S5: **Distribution of optimized muscle parameters.** Final values for maximum contraction velocity (top), optimal fiber length (middle), and maximum isometric force (bottom) after optimization. Each dot represents one of the top 10 individuals selected from a total population of 200. Muscles are grouped by segment (thorax, coxa, femur) from left to right.

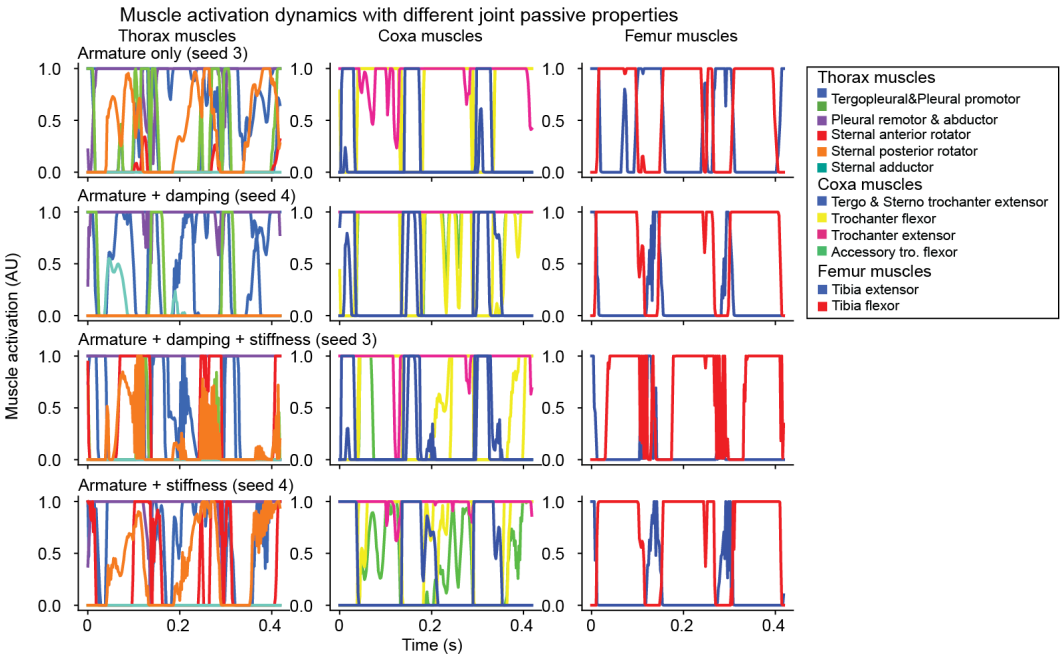

Figure S6: **Muscle activation dynamics across passive joint conditions.** Simulated activations of thorax (left), coxa (middle), and femur (right) muscles in PPO agents trained with different joint properties: (i) armature only (top), (ii) armature and damping (second), (iii) armature and damping and stiffness (third), and (iv) armature and stiffness (bottom). Shown are rollouts from the best-performing seed in each condition.

REFERENCES

Moritz Haustein, Alexander Blanke, Till Bockemühl, and Ansgar Büschges. A leg model based on anatomical landmarks to study 3d joint kinematics of walking in drosophila melanogaster. *Frontiers in Bioengineering and Biotechnology*, 12:1357598, 2024.

Victor Lobato-Rios, Shravan Tata Ramalingasetty, Pembe Gizem Özdil, Jonathan Arreguit, Auke Jan Ijspeert, and Pavan Ramdya. NeuroMechFly, a neuromechanical model of adult Drosophila melanogaster. *Nature Methods*, 19(5):620–627, May 2022. ISSN 1548-7105. doi: 10.1038/s41592-022-01466-7.

Pembe Gizem Özdil, Auke Ijspeert, and Pavan Ramdya. SeqIKPy: a python package for inverse kinematics in insects, 2024. URL `https://doi.org/10.5281/zenodo.12601317`.

Arnault H Caillet, Andrew TM Phillips, Christopher Carty, Dario Farina, and Luca Modenese. Hill-type computational models of muscle-tendon actuators: a systematic review. *bioRxiv*, pages 2022–10, 2022.

Örjan Ekeberg, Marcus Blümel, and Ansgar Büschges. Dynamic simulation of insect walking. *Arthropod Structure & Development*, 33(3):287–300, July 2004. ISSN 1467-8039. doi: 10.1016/j.asd.2004.05.002.

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.

Hartmut Geyer, Andre Seyfarth, and Reinhard Blickhan. Positive force feedback in bouncing gaits? *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 270(1529):2173–2183, 2003.

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.

Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE transactions on evolutionary computation*, 6(2): 182–197, 2002.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017.

Vittorio La Barbera, Fabio Pardo, Yuval Tassa, Monica Daley, Christopher Richards, Petar Kormushev, and John Hutchinson. Ostrichrl: A musculoskeletal ostrich simulation to study biomechanical locomotion. *arXiv preprint arXiv:2112.06061*, 2021.