Computational Design of Terrestrial Robots with Anisotropic Friction

HANG HU, Tsinghua University, China and Shanghai Qi Zhi Institute, China
KANGBO LYU, Shanghai Qi Zhi Institute, China and Tsinghua University, China
CHANGYU HU, Tsinghua University, China
ZIHAN LI, Tsinghua University, China
PEIWEN YANG, Tsinghua University, China
MINCHEN LI, Carnegie Mellon University, United States of America and Genesis AI, United States of America
SHUGUANG LI, Tsinghua University, China
TAO DU, Tsinghua University, China and Shanghai Qi Zhi Institute, China

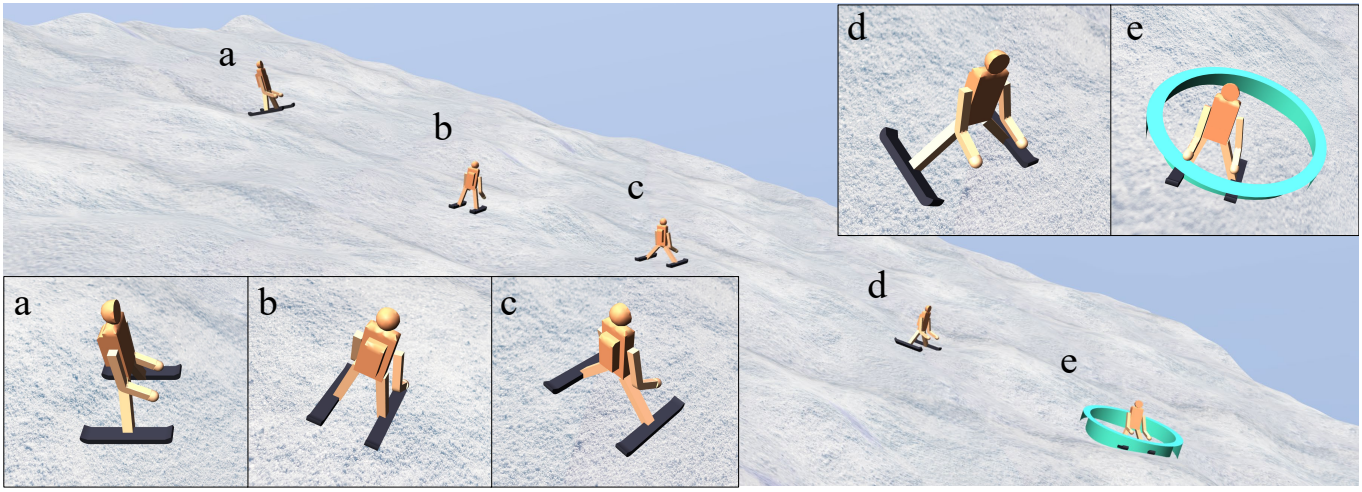

Fig. 1. A robot skis down a slanted terrain by jointly optimizing its actuator control policy and the orientation of anisotropic friction on its ski boards. The two ski boards adopt distinct optimized friction orientations of 19° and 64°, defined as the angle between the forward direction and the principal axis of the anisotropic friction configuration. As shown in (a–e), the robot performs smooth curved trajectories, sharp turns, and a sudden braking maneuver to stop precisely at the target, showcasing the agile locomotion enabled by the anisotropic friction and our proposed pipeline.

Anisotropic friction is a critical source of propulsion for efficient locomotion of many terrestrial animals. The interplay between animal morphology, control, and anisotropically frictional contact makes designing optimal anisotropic friction for terrestrial locomotion intriguing and challenging. We propose a computational pipeline for co-designing anisotropic friction and controllers of terrestrial robots with diverse morphologies. Our pipeline presents a co-design algorithm that alternates between optimizing direction of anisotropic friction and training a neural network controller to improve the locomotion performance of a given robot morphology. Based on the intuition that controller's performance does not change significantly when the frictional force differs slightly, we introduce the concept of trust-region into robot co-design, allowing the controller network to continue training from the previous iteration. Our evaluation on various morphologies show that anisotropic friction is critical for terrestrial robot locomotion, and our pipeline is statistically better than current state-of-the-art methods. Furthermore, we reveal that large language models (LLM) constitute a strong baseline for this kind of co-design problems, worth receiving more attention. We demonstrate that co-designing anisotropic friction and control unlocks effective locomotion in various downstream tasks, including locomotion on uneven terrain, navigation in a maze, and object manipulation. To validate our pipeline in the real world, we design and 3D print a variety of scales and systematically measure their anisotropic friction coefficients. Then we construct a multi-link robot with anisotropic scales designed by our pipeline and compare its performance with isotropic scales. Our real-world experiments confirm that isotropic scales are insufficient to support terrestrial robots' locomotion abilities, and computationally co-designing friction and control enables robots to perform tasks including turning, slithering, and other non-trivial locomotion tasks.

CCS Concepts: • **Computing methodologies** → **Modeling and simulation**.

Authors' Contact Information: Hang Hu, Tsinghua University, Beijing, China and Shanghai Qi Zhi Institute, Shanghai, China, huh23@mails.tsinghua.edu.cn; Kangbo Lyu, Shanghai Qi Zhi Institute, Shanghai, China and Tsinghua University, Beijing, China; Changyu Hu, Tsinghua University, Beijing, China; Zihan Li, Tsinghua University, Beijing, China; Peiwen Yang, Tsinghua University, Beijing, China; Minchen Li, Carnegie Mellon University, Pittsburgh, United States of America and Genesis AI, Pittsburgh, United States of America; Shuguang Li, Tsinghua University, Beijing, China; Tao Du, Tsinghua University, Beijing, China and Shanghai Qi Zhi Institute, Shanghai, China.

SIGGRAPH Conference Papers '26, July 19–23, 2026, Los Angeles, CA, USA

ACM ISBN 979-8-4007-2554-8/2026/07
https://doi.org/10.1145/3799902.3811103

Additional Key Words and Phrases: computational design and fabrication; biomimetic robots

ACM Reference Format:
Hang Hu, Kangbo Lyu, Changyu Hu, Zihan Li, Peiwen Yang, Minchen Li, Shuguang Li, and Tao Du. 2026. Computational Design of Terrestrial Robots with Anisotropic Friction. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '26), July 19–23, 2026, Los Angeles, CA, USA.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3799902.3811103

1 Introduction

Many terrestrial animals, such as snakes, anguid lizards, and amphisbaenians, leverage anisotropic friction for locomotion [Gans 1962; Hu et al. 2009; Marvi et al. 2014]. The significant role of anisotropic friction in locomotion, as opposed to its isotropic counterpart, can be intuitively demonstrated using simple computational examples (Fig. 2) and has also been confirmed by previous work [Qi et al. 2023]. The efficacy of anisotropic friction in locomotion has inspired researchers to propose designs of biomimetic robots, e.g., snake robots equipped with scales that exhibit anisotropic friction for slithering [Marvi et al. 2011; Rafsanjani et al. 2018; Serrano et al. 2015]. Compared with these manual robot designs, computationally designing robots that exploit friction for locomotion is much less explored due to the intricate interplay between their friction and control, despite the promising success of computational robot design in related topics [Hu et al. 2023; Jeon et al. 2025; Lu et al. 2025].

Our work studies the computational design of terrestrial robots that exploit anisotropic friction to accomplish locomotion tasks. Unlike many existing robot co-design works that explore robot link and joint designs and their controllers [Gupta et al. 2022, 2021; Hu et al. 2023; Jeon et al. 2025; Lu et al. 2025; Zhao et al. 2020], we focus on a much less explored topic: co-optimizing anisotropic friction and control for a given robot link and joint design. To achieve this goal, we develop a computational pipeline that models an articulated rigid-body robot with anisotropic frictional contact, simulates its motion under closed-loop control, co-optimizes its anisotropic friction and control, and finally, creates its fabrication plan for sim-to-real deployment (Fig. 3).

At the core of our pipeline is a novel co-optimization algorithm to address the key challenge in *bi-level* optimization typically encountered in robot co-design: evaluating a robot design (in our case, the anisotropic friction in each link) requires first solving an optimal controller for it, which typically involves expensive reinforcement learning (RL). To tackle this challenge, we draw inspiration from the *trust-region* method [Nocedal and Wright 2006] and develop a new robot co-optimization algorithm with two key insights: First, we update our anisotropic friction design within a trust region and maintain a controller shared by all designs in the trust region, eliminating repeated RL training whenever exploring a new design. Second, we update the shared controller whenever the friction design and its trust region are updated. Because the trust region allows only mild changes to the friction design, we have observed that it is typically sufficient to run modest fine-tuning rather than full RL training to update the shared controller. These two insights lead to a novel and effective co-optimization algorithm for our problem.

We compare our co-optimization algorithm with competitive baselines to analyze its performance. We identify BodyGen [Lu et al. 2025], one of the latest methods in robot co-design, as a state-of-the-art (SOTA) baseline. We also construct robot co-designs driven by GPT-5 and Qwen3, two recent large language models (LLMs), as baselines. Such LLM-driven designs have not been fully explored in robot co-design research. Yet, their performance in our problem is highly competitive compared with BodyGen, suggesting the potential of LLMs for robot co-design. Compared with these baselines, our co-optimization algorithm consistently discovers designs with substantially higher performance on representative robot morphologies after repeated experiments, highlighting the value of introducing trust-region methods to robot co-design.

We demonstrate the efficacy of our computational design pipeline by co-optimizing a robot's anisotropic friction and control to accomplish locomotion tasks that require substantial usage of anisotropic friction, including navigating in a maze (Fig. 7), delivering an object (Fig. 8), and skiing with a hard stop (Fig. 1). Finally, we transfer the optimized anisotropic friction design and control to a real-world robot and demonstrate its effective turning, slithering, and point-tracking capability. Notably, some of these tasks are known to be challenging for robots with isotropic friction, as confirmed by our hardware experiments (Sec. 8.2) and prior works [Qi et al. 2023]. Code and hardware data for this paper are at https://github.com/Hang-Hu1/snake.git.

We summarize our contributions below:

(1) We develop a computational design pipeline for simulating, co-optimizing, and fabricating robots that leverage anisotropic friction to accomplish locomotion tasks;
(2) We propose a robot co-optimization algorithm that presents a novel adoption of trust-region methods within a classic bi-level optimization scheme for robot co-design problems;
(3) We conduct extensive evaluations to demonstrate that, for our problem, our robot co-optimization algorithm achieves substantially better performance than competitive baselines, including several variants of the state-of-the-art (SOTA) robot co-design methods and designs driven by the large language models (LLMs).
(4) We demonstrate the efficacy of our pipeline by co-designing anisotropic friction and robot control to accomplish various tasks in simulation and in the real world.

2 Related Work

Computational design of robots. Computational design of robots explores computational methods that jointly optimize robot morphology (e.g., robot shape or link-joint connection) and control. Previous research in this field has studied the design of a wide range of robots, including rigid and soft robots in flying, walking, and swimming tasks [Du et al. 2016; Gupta et al. 2022, 2021; Hu et al. 2023; Le et al. 2024; Lu et al. 2025; Ma et al. 2021; Umetani et al. 2014; Wang et al. 2023; Zhao et al. 2020]. Our work falls into this category but differs from existing research by studying robot co-design problems involving anisotropic friction and control for locomotion tasks, a topic unexplored in many computational design papers.

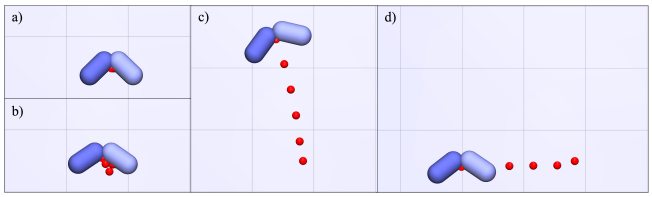

Fig. 2. *Motivating Example.* A planar two-link robot is actuated by a periodic signal at the hinge joint. The trajectory of the robot's center of mass is illustrated by red dots. **(a)** isotropic friction $\mu = 0.2$; **(b)** isotropic friction $\mu = 0.8$; **(c)** anisotropic directional friction with $\mu_+ = 0.2$ and $\mu_- = 1.0$; **(d)** anisotropic left-right friction with $\mu_{\text{left}} = 0.2$ and $\mu_{\text{right}} = 1.0$. Despite identical actuation, the robot exhibits negligible net displacement under isotropic friction ((a) and (b)), whereas anisotropic friction induces smooth and directed locomotion ((c) and (d)).

Friction in locomotion. Many terrestrial animals utilize anisotropic friction to generate efficient locomotion gaits. One of the most well-known and extensively studied examples is snake slithering [Hu et al. 2009; Marvi et al. 2013; Marvi and Hu 2012; Marvi et al. 2011]. A large number of studies in mechanical engineering have built real-world robots to reproduce the anisotropically frictional locomotion [Han et al. 2015; Ji et al. 2022; Qi et al. 2023; Serrano et al. 2015; Shen et al. 2021; Tramsen et al. 2018]. Previous work [Qi et al. 2023] has also demonstrated that a biomimetic snake robot with *isotropically frictional scales* struggles to move effectively, confirming the significance of studying the design of anisotropic friction in such locomotion problems. We differ from the biomimetic research in this category in that we contribute a computational pipeline for exploring unseen designs, whereas biomimetic research emulates known designs in nature.

Optimization methods. Many previous studies treat robot co-design as a bi-level optimization problem: a high-level optimizer selects a robot morphology, and a low-level optimizer finds an effective controller for that morphology. Optimizers for selecting robot morphologies include Bayesian optimization (BO) [Hu et al. 2023], graph heuristic search [Zhao et al. 2020], evolutionary algorithms [Gupta et al. 2021], and large language models (LLMs) [Song et al. 2025]. Popular optimizers for the low-level control include model predictive control [Hu et al. 2023; Zhao et al. 2020] and reinforcement learning (RL) [Gupta et al. 2021]. Similar to these studies, our work combines BO with RL, and our novelty lies in proposing a trust-region (TR) method that bridges the morphology and control optimizers in this bi-level optimization of robots.

Unlike the bi-level optimization discussed above, BodyGen [Lu et al. 2025] and other recent robot co-design studies [Gupta et al. 2022; Yuan et al. 2022] advocate for jointly solving shape and control in a single, unified optimization problem. We have investigated this line of research and found that, while solving it yields effective co-design solutions for our problem setup, its locomotion performance is substantially worse than that of our bi-level optimization method or a preliminary LLMs-driven design.

3 Design Space

This work considers rigid robots with articulated joints in contact-rich locomotion tasks. While co-designing morphology and brain for such locomotion robots is not a new problem [Gupta et al. 2021; Hu et al. 2023; Lu et al. 2025; Zhao et al. 2020], morphology design in existing works only explores structures (i.e., different link-joint connections and shapes of links) and does not consider the influence of different friction profiles, especially the design of anisotropic friction. Our research differs from these works and is complementary, as we focus on co-designing the anisotropic friction of each link and the controller for each joint, given fixed robot structures.

Robot topology. The topology of a robot is defined as an undirected graph $G = (V, E)$, where each node is a link of a robot, and each edge represents a hinge joint connecting two links.

Anisotropic friction design. We use an 8-point radar chart to describe the friction profile of a link, where the 8 values represent the friction coefficients at 8 points spaced 45° apart on the contact plane (Fig. 3), which appears in bio-inspired design of robots [Qi et al. 2023]. With this parametrized friction profile, we adopt the classic maximal dissipation principle [Andrews et al. 2022] to compute the frictional force.

Based on our observations from extensive simulation and hardware experiments, we find that given an anisotropic friction profile, changing the direction of the profile (rotating the profile in the radar chart) has a more significant impact on the locomotion performance than perturbing the 8 components of the profile. Taking the simplest 2-link robot in Fig. 2 as an example, when the friction profiles of links are rotated by a certain angle, the robot's trajectory will shift by approximately the same angle. Therefore, in order to increase optimization efficiency without reducing too much representation capability, we fix the friction profile for each link and instead optimize its orientation. From this idea, each link $V_i \in V$ is assigned a real number $d_i \in [0, 2\pi)$ to denote its orientation. The design space of anisotropic friction is the set of orientations $\boldsymbol{d} = \{d_i | V_i \in V\}$.

Controller design. We apply a standard multi-layer perceptron (MLP) neural network controller to our randomly generated robot morphology with scales:

$$\boldsymbol{a} = \pi_{\boldsymbol{w}}(\boldsymbol{s}), \tag{1}$$

where π is the network controller with $\boldsymbol{w}$ its learnable parameter, $\boldsymbol{s}$ the state of the robot, and $\boldsymbol{a}$ the generated action. We define the robot's state as $\boldsymbol{s} = (\boldsymbol{d}, \boldsymbol{q}, \dot{\boldsymbol{q}}, \boldsymbol{k})$ where $\boldsymbol{d}$ denotes the orientations of the links, $\boldsymbol{q}$ represents the generalized coordinates formed by the joint angles, $\dot{\boldsymbol{q}}$ denotes the corresponding generalized velocities, and $\boldsymbol{k}$ denotes auxiliary information related to the task. The action $\boldsymbol{a}$ comprises the reference position of each hinge joint, which is then used to compute the torque via a PID controller.

4 Simulation

The design process in Sec. 3 generates robots with rigid links, articulated joints, and anisotropically frictional scales. Computational design of such robots in contact-rich locomotion tasks requires a simulator which carefully balances among speed, physical accuracy, and robustness to support extensive evaluations of massive robot

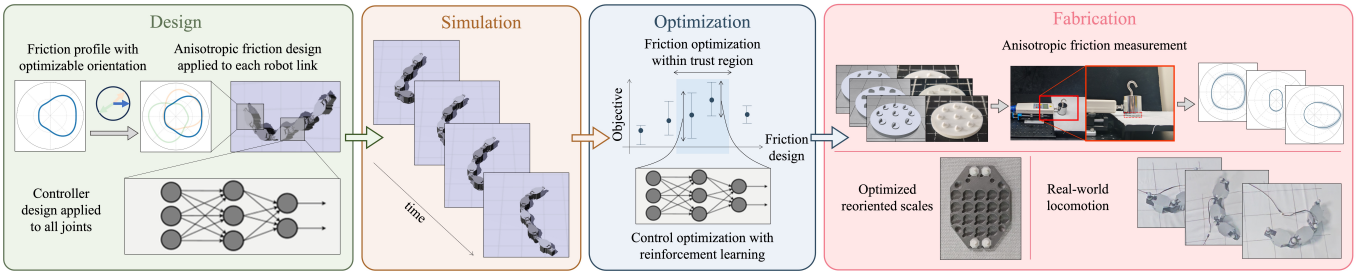

Fig. 3. *Pipeline Overview.* Our pipeline consists of four stages. *Design:* we co-design the friction profile by optimizing the orientations of anisotropic scales on each robot link, together with the controller parameters applied to actuated joints (Sec. 3). *Simulation:* we develop a customized physics-based simulator to evaluate locomotion performance (Sec. 4). *Optimization:* a two-level optimization framework alternately performs friction profile optimization using a trust region strategy and controller optimization with reinforcement learning (Sec. 5). *Fabrication:* we measure anisotropic friction properties of 55 3D-printed scales, fabricate the robot with optimized reoriented scales, and deploy the learned closed-loop control policies for real-world locomotion experiments (Sec. 6).

designs. For these reasons, we use MuJoCo [Todorov et al. 2012], a classic simulator widely adopted in robot learning research [Brockman et al. 2016; Lu et al. 2025; Tassa et al. 2018], as our backbone simulator. Because the current version of MuJoCo (3.3.0) does not support our definition of friction profiles, we extend it with our customized implementation of anisotropic friction (Alg. 1).

Equations of motion. The forward dynamics of our robots are formulated as a Linear Complementarity Problem (LCP) subject to contact constraints [Andrews et al. 2022]. The discretized equations of motion can be generally written as:

$$\mathbf{M}(\boldsymbol{q})\ddot{\boldsymbol{q}} + \boldsymbol{c}(\boldsymbol{q}, \dot{\boldsymbol{q}}) = \boldsymbol{\tau} + \mathbf{J}^T(\boldsymbol{f}_n + \boldsymbol{f}_c) + \boldsymbol{f}_{ext}, p \tag{2}$$

where $\mathbf{M}$, $\boldsymbol{c}$, and $\boldsymbol{\tau}$ represent the mass matrix, nonlinear forces, and actuation torques, respectively. The terms $\mathbf{J}^T\boldsymbol{f}_n$ and $\mathbf{J}^T\boldsymbol{f}_c$ handles the normal force and friction force in contact interactions, respectively, with $\mathbf{J}$ being the constraint Jacobian[Andrews et al. 2022]. Crucially, we incorporate anisotropic friction into the friction force $\boldsymbol{f}_c$ using the *maximal dissipation principle* [Andrews et al. 2022]. Given the sliding velocity $\boldsymbol{v}_c$, the orientation-dependent friction coefficient $\mathcal{M}$ (Sec. 3), the directions of friction profiles $\boldsymbol{d}$ and the normal force f_n, the frictional force $\boldsymbol{f}_c$ is solved by maximizing the energy dissipation:

$$\boldsymbol{f}_c = \underset{\boldsymbol{f}\in\mathcal{F}(\mathcal{M},\boldsymbol{d},f_n)}{\arg\min} \; \boldsymbol{v}_c \cdot \boldsymbol{f}, \tag{3}$$

ensuring the friction force optimally resists motion within the feasible region $\mathcal{F}$ defined by our anisotropic scales.

Time-stepping with anisotropic friction. We solve Eq. 2 with the semi-implicit Euler time integration provided in MuJoCo. We inherit the majority of computation in MuJoCo's semi-implicit Euler implementation but override its frictional force with our implementation of anisotropic friction (Alg. 1): assuming a link with its anisotropic friction profile is in contact, we first compute its contact velocity $\boldsymbol{v}_c^{(i)}$ and constraint force $\boldsymbol{f}_n^{(i)}$ (line 6). Then, we compute the anisotropic friction $\boldsymbol{f}_c^{(i)}$ using the maximal dissipation principle (line 8), and add it as an explicit force into the dynamic system (line 10).

Validation of anisotropic friction. To validate our customized MuJoCo module, we reproduce a counterintuitive phenomenon of the anisotropic friction [Erleben et al. 2020]. Unlike isotropic friction, which is typically modeled as acting opposite to the direction of relative velocity [Li et al. 2020], anisotropic frictional forces need not be aligned with the relative velocity [Andrews et al. 2022] and can lead to counter-intuitive curved trajectories. As shown in Fig. 4, boxes subjected to impulses of equal magnitude but different directions exhibit trajectories that bend toward low-friction directions, confirming that the implemented module behaves as expected.

Algorithm 1 Simulation Within One Timestep

1: **Input:** State of the m-th timestep $\boldsymbol{q}_m$, friction profiles $\mathcal{M}$, link directions $\boldsymbol{d}$
2: **Output:** State of the $(m+1)$-th timestep $\boldsymbol{q}_{m+1}$
3: $C \leftarrow$ MuJoCo.CollisionDetection($\boldsymbol{q}_m$)
4: $\boldsymbol{F}_{ext} \leftarrow []$
5: **for** each contact $i \in C$ **do**
6: $\quad \boldsymbol{v}_c^{(i)}, \boldsymbol{f}_n^{(i)} \leftarrow$ SolveContactVelocityAndConstraintForce(i)
7: $\quad \mu^{(i)} \leftarrow$ FindFrictionProfile($\mathcal{M}, i$)
8: $\quad$ Solve friction $\boldsymbol{f}_c^{(i)} \leftarrow \arg\min_{\boldsymbol{f}\in\mathcal{F}(\mu^{(i)},\boldsymbol{d}^{(i)},\boldsymbol{f}_n^{(i)})}\{\boldsymbol{v}_c^{(i)} \cdot \boldsymbol{f}\}$
9: $\quad \boldsymbol{F}_{ext} \leftarrow \boldsymbol{F}_{ext}.append(\boldsymbol{f}_c^{(i)})$
10: **end for**
11: $\boldsymbol{q}_{m+1} \leftarrow$ MuJoCo.SemiImplicitEulerIntegration($\boldsymbol{q}_m, \boldsymbol{F}_{ext}$)
12: **return** $\boldsymbol{q}_{m+1}$

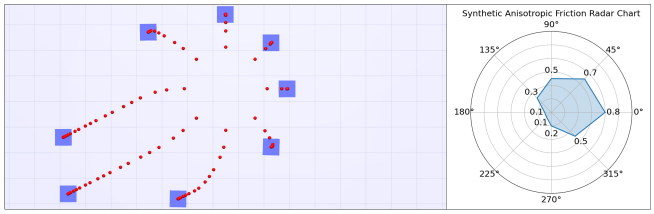

Fig. 4. Validation of anisotropic friction in our customized MuJoCo simulator. **Left:** 8 boxes with identical anisotropic friction parameters are initialized on a unit circle and pushed by outward-pointing impulses. The resulting curved trajectories demonstrate the frictional force, determined by the maximum dissipation rule, may not align with the velocity direction. **Right:** The directional friction coefficients used, visualized by the radar chart.

5 Optimization

The optimization problem. With the definitions in Sec. 3, we now provide a formal definition of our optimization problem. Given the topology of a robot $G = (V, E)$, where $V = \{V_1, \ldots, V_k\}$, the friction profile of each link $\mathcal{M} = \{\mu_1, \ldots, \mu_k\}$, and the task reward function $f(\cdot)$, We aim to optimize the orientation of each link $\boldsymbol{d} = \{d_1, \ldots, d_k\}$ and the controller π to maximize the reward. Specifically, the problem is formulated as:

$$\boldsymbol{d}^* = \arg\max_{\boldsymbol{d}} f(\boldsymbol{d}, \pi^*_{\boldsymbol{d}}) \tag{4}$$

$$s.t.\ \ \pi^*_{\boldsymbol{d}} = \arg\max_{\pi_{\boldsymbol{d}}} f(\boldsymbol{d}, \pi_{\boldsymbol{d}}). \tag{5}$$

This problem has some key properties. First, it is a co-optimization problem because it couples both friction and control: the friction parameters alter the robot's dynamics, and conversely, the controller's optimization determines the quality of the selected friction parameters. Second, it is a bilevel problem because the control optimization is nested within the friction optimization loop. Third, the degrees of freedom (DoFs) in the friction and control optimization are highly imbalanced: the number of links $|\boldsymbol{d}|$ is typically small (e.g., < 20), whereas the controller parameter involves a vastly larger search space, often exceeding 20,000 dimensions.

Challenges. Robot co-design of morphology (topology, shape, or in our case, friction) and the controller is a classic problem in computational design [Gupta et al. 2021; Hu et al. 2023; Jeon et al. 2025; Lu et al. 2025; Zhao et al. 2020]. Many studies have proposed different combinations of optimizers for morphology and control to tackle this bilevel optimization problem, e.g., combine Covariance Matrix Adaptation Evolution Strategy (CMA-ES) for the high-level shape optimization with RL for the low-level control optimization [Gupta et al. 2021], or combine Bayesian Optimization (BO) or graph heuristic search (GHS) for shape optimization and model predictive control (MPC) for control optimization [Hu et al. 2023; Zhao et al. 2020]. However, the bilevel optimization results in extensive time spent running RL to optimize the low-level controller, while MPC sacrifices the power of neural network control unlocked by state-of-the-art robot learning. An alternative is to train a universal controller [Gupta et al. 2022] to accommodate different morphology parameters, which requires a collection of diverse morphology examples beforehand and expensive pretraining on it.

Our method: trust-region robot co-design. The core bottleneck in the aforementioned bilevel optimization lies in the evaluation step: to strictly score a single friction design candidate, one must train a corresponding controller from scratch using RL, which is computationally prohibitive to repeat within an outer optimization loop. A natural strategy to amortize this cost is transfer learning—initializing the controller for a new design using the converged policy from a previous design. This allows for rapid fine-tuning rather than full training. However, the effectiveness of such transfer relies heavily on the similarity between the two tasks. If the friction parameters change too drastically (e.g., flipping the anisotropic direction), the previous controller may be ill-suited or even detrimental to the new environment, causing fine-tuning to fail. Based on this intuition, we argue that to fully leverage the efficiency of controller reuse, we must constrain the search process to ensure that consecutive friction designs remain structurally similar. This naturally motivates us to adopt a trust-region approach [Nocedal and Wright 2006]. By confining the friction parameter updates within a dynamic local region, we ensure that the dynamics change smoothly, thereby guaranteeing that the previous controller serves as a valid "warm start" for the new design. This insight leads to our proposed Trust-Region Robot Co-design algorithm (Alg. 2).

Specifically, we initialize a friction parameter and controller and gradually update them throughout co-optimization. In each iteration, we employ a trust-region update to the friction parameter: given the current friction parameters, we sample candidate designs from its neighborhood utilizing the historical data. Next, we perform PPO updates to fine-tune the controller using motion data generated with the new friction parameters. Upon completion, the new design-controller pair is scored to determine whether to expand or shrink the trust region. The trust region is always centered at the best friction parameter so far. If the number of consecutive successes reaches δ_{succ}, the side length of the trust region is doubled; if the number of consecutive failures reaches δ_{fail}, the side length of the trust region is halved. Each co-optimization is repeated n times.

We now explain the insights in our co-design algorithm. First, the trust region constraint serves as a guarantee for efficient adaptation. By confining the friction parameters to the vicinity of the current design, the previous controller acts as a reliable warm-start, requiring only a few PPO iterations to adapt to the mild dynamics changes. Second, the trust region mechanism provides robustness against suboptimal controller updates. In cases where the limited PPO updates fail to yield a satisfactory controller (indicating a large "sim-to-sim" gap), the resulting low performance score triggers a reduction in the trust region size. This shrinkage implicitly forces the algorithm to explore a smaller neighborhood in subsequent iterations, thereby allowing the controller more opportunities to converge within the local dynamics space. The complete procedure is detailed in Alg. 2.

Algorithm 2 Co-Design Pipeline

Input: Topology $G = (V, E)$, friction profiles $\mathcal{M}$
Output: Optimized friction parameters $\boldsymbol{d}$ and controller π
Initialize decision variables $\boldsymbol{d}_0$ and $\boldsymbol{w}_0$ randomly
$d_{hist} \leftarrow [], f_{hist} \leftarrow []$
for $i = 1$ to n **do**
  $\boldsymbol{d}_i \leftarrow \text{TrustRegionUpdate}(\pi_{\boldsymbol{w}_{i-1}}, d_{hist}, f_{hist})$
  $\boldsymbol{w}_i \leftarrow \text{UpdateController}(\boldsymbol{d}_i, \boldsymbol{w}_{i-1})$
  Scoring: $y \leftarrow f(\boldsymbol{G}, \mathcal{M}, \boldsymbol{d}_i, \pi_{\boldsymbol{w}_i})$
  $d_{hist} \leftarrow d_{hist}.append(\boldsymbol{d}_i), f_{hist} \leftarrow f_{hist}.append(y)$
end for
return $\boldsymbol{d}_n, \pi_{\boldsymbol{w}_n}$

Remarks. While our algorithm draws inspiration from standard Trust-Region Bayesian Optimization (TRBO) [Eriksson et al. 2019], it introduces a critical adaptation tailored for the bilevel nature of robot co-design. BO applied to the friction parameters treats control optimization as a black box, spending extensive time optimizing control for each individual shape in a BO sample [Konakovic Lukovic

et al. 2020]. A vanilla TRBO applied to our problem would extend BO by introducing TR to guide updates to friction parameters, but would still leave control optimization within the evaluation of each BO sample. We take a step further by separating the controller optimization from the BO samples and treating the controller as part of the trust-region model, which is gradually updated throughout optimization. Empirical evidence shows that our algorithmic design is simple yet powerful, substantially improving BO and outperforming bilevel optimization baselines, including those driven by large language models (LLMs) (Sec. 7).

6 Fabrication

We propose a modular design for each link. A link consists of a top cover, an upper panel, a body, a bottom panel, and a chassis, which are assembled together in sequence (Fig. 5). The upper and bottom panels are responsible for connecting the link with at most two servos at the two ends. The body is a hollow shell that allows additional weights to be inserted. The chassis is designed with many small holes for installing scales.

We use hinge joints to connect these links. Each joint is implemented using a dual-axis positional servo (Hiwonder LX-16A) that connects two links. The axis of rotation of each joint is parallel to the vertical axis. We use a computer or a servo controller to provide signals indicating their reference positions.

Finally, we design and measure the anisotropicities of 55 3D-printed scales, and select a representative one to deploy on our robots throughout the hardware examples (Sec. 8.2).

7 Evaluation

Tasks. We design three robots whose topologies are similar to the letters "S", "I", and "Y" as our evaluation benchmark (Fig. 6). We equip each link in these robots a friction profile whose eight frictional coefficients are selected from a binary set $\{0.2, 1.0\}$, which mimics the scales observed on terrestrial animals. We consider a point navigation task for both robots: We randomly select a direction pointing from the head link of a robot, create a target about one body length away, and repeat this process whenever the robot reaches the target. We use 1 ms as our simulator's step size and run the control signal at 2 Hz for this task.

Baselines. As no existing robot co-design pipelines are readily available for optimizing anisotropic friction and control, we use SOTA robot co-design algorithms from similar settings and adapt them to our problem. These include 1) BodyGen [Lu et al. 2025], one of the latest studies for co-designing terrestrial robots for locomotion. We replace its morphology design variables with our anisotropic friction design variables and update its network inputs and outputs accordingly; 2) BodyGen-MLP, which replaces the Transformer-based network in BodyGen with the same multi-layer perceptron (MLP) network as in our algorithm; 3) BO + RL, a vanilla co-design algorithm that uses BO to select anisotropic friction design and trains RL controllers to evaluate each design; 4) LLM-GPT-5 and LLM-Qwen3, which replace the friction design in our algorithm with a descrption of the current problem context in a prompt sent to LLMs and ask for adjustments on the friction design in return.

Training. We train all methods with 10 million control steps, i.e., 5 billion simulation steps. We use the same hyperparameters for our method and all baselines, except for BodyGen, which has a different transform-based network architecture and requires additional hyperparameter fine-tuning. We repeat each method with three random seeds on each robot.

Results and discussions. Fig. 6 summarizes the average locomotion speed and its standard deviation over successful completions of point navigation using all methods as their training progresses on the three robots. We refer readers to our video for the detailed locomotion results. Our method achieves competitive performance across all baselines, sometimes with substantial gains, as indicated by its higher locomotion speed during training. In particular, the performance boost over vanilla BO confirms the efficacy of integrating the trust-region method in our robot co-design pipeline.

Comparisons between LLM and non-LLM baselines reveal an insightful understanding of them. Our use of both LLMs is largely brute-force. Yet, they achieve competitive performance in our problem, sometimes surpassing the SOTA robot co-design methods. This observation implies that robot co-design research should consider LLMs as strong baselines with promising potential.

Finally, we observe that, regardless of whether the methods are LLM-based, no single method consistently dominates the others across the three robots. The observation indicates the complexity of this robot co-optimization problem. One factor behind this observation is that the optimal locomotion capabilities of these three robots may be inherently determined by their letter topologies, which could favor different co-optimization algorithms.

8 Examples

8.1 Examples in Simulation

Navigating a maze. This example deploys the I-robot (Sec. 7) in a 2D grid-based maze. We run a breadth-first search (BFS) in the maze map to generate a sequence of targets, which we send to the point-navigation controller trained by our method as explained in Sec. 7. The I-robot demonstrates smooth locomotion between the starting and ending points in the maze, including slithering for straight-line forward motion and occasional sharp U-turns enabled by the optimized anisotropic friction and control (Fig. 7).

Object manipulation. This example co-designs the S-robot (Sec. 7) for challenging nonprehensile object manipulation, where the robot must propel an object to a target using only body undulation without grasping end-effectors. We select five objects with distinct geometries (*capsule*, *cube*, *dumbbell*, *T-shape*, and *tetris-Z*) and initialize them at random locations. The objective combines distance-based costs and a completion bonus. The optimization consistently converges within 8 hours on a standard workstation.

Our optimized designs successfully complete this task for all objects (Fig. 8), significantly outperforming random initializations which typically fail to generate sufficient traction or slip sideways. Our optimized locomotion does not converge on a single fixed gait but instead discovers diverse manipulation strategies adapted to the immediate state. For *capsule* and *cube*, the robot exhibits and switches flexibly between"caging" (curling the body to constrain

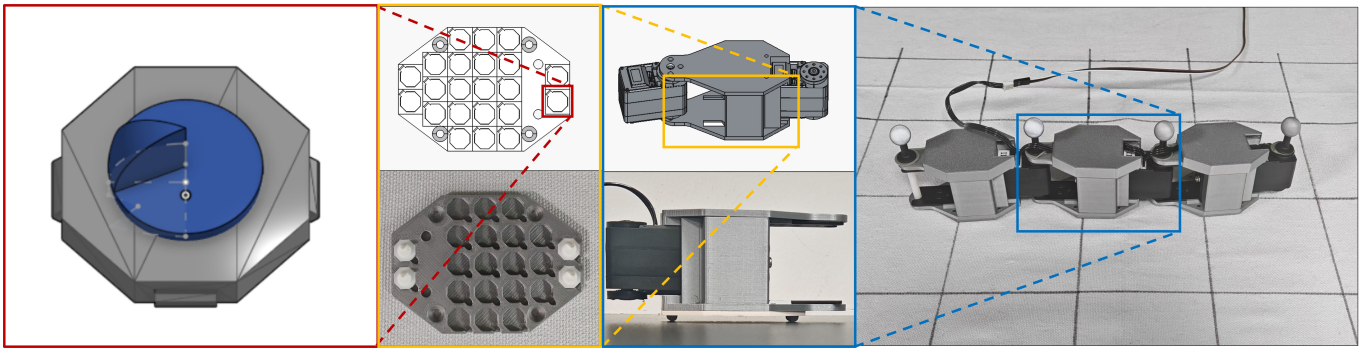

Fig. 5. Illustration of the fabrication process of our pipeline. The first column demonstrates the details of the scale we select. The second column shows the CAD model and a real photo of a chassis that uses a buckle to deploy scales. The third column gives the CAD model and a real photo of the link. This is a modular design that can be repeated multiple times. The last column illustrates our robot's full view.

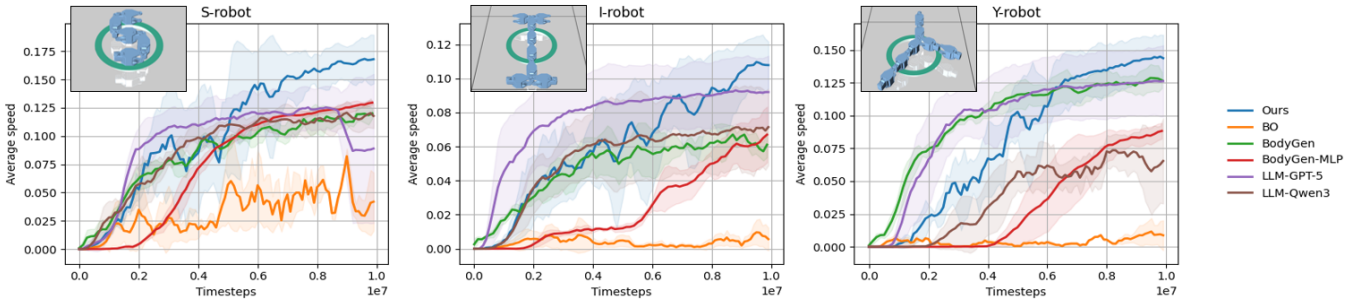

Fig. 6. Results of our comparison experiment. Each plot demonstrates the evaluation result on one robot (with the robot structure at the top left). Each curve represents the task completion speed in the process of training. The shaded region surrounding the curve indicates the standard deviation of this method. Results are averaged over 3 to 5 runs with different random seeds. It can be concluded that our method surpasses all the baselines.

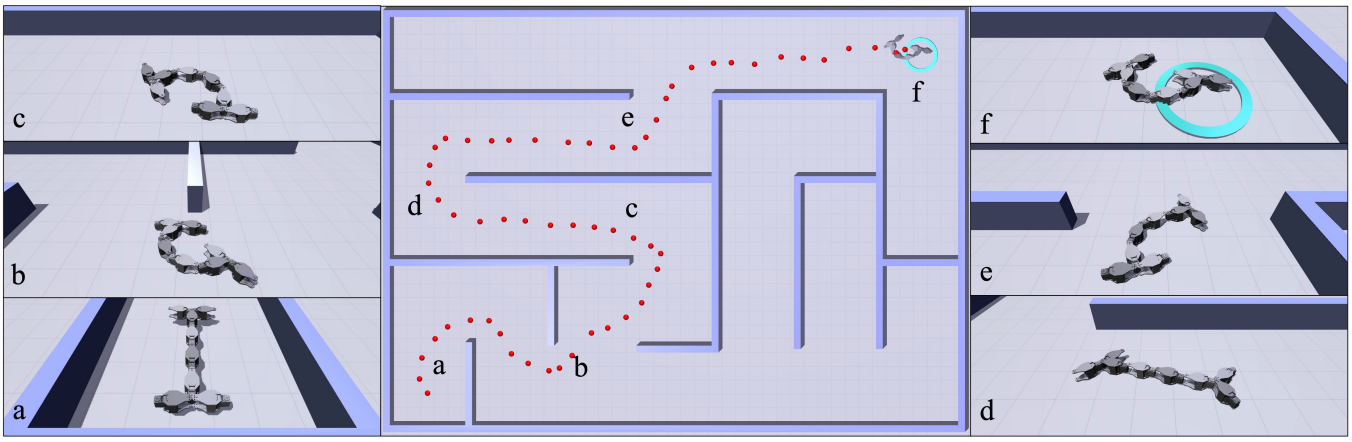

Fig. 7. The *I-robot* navigating a maze. *Middle*: Overview of the 2D grid-based maze navigation task, with the center of mass trajectory shown as red dots. *(a-f)*: Representative snapshots taken at the corresponding points along the trajectory, illustrating the robot's configurations and locomotion modes during maze navigation, including straight-line forward motion and turning maneuvers.

the object) and "pushing" behaviors (using the head link to propel the object). For *dumbbell* and *tetris-Z* with irregular shapes, the robot aligns its body segments parallel to the object's flat face to maximize its contact area. Crucially, the optimized friction orientations exhibit functional heterogeneity: the robot learns to orient certain links to minimize forward drag and use others with high friction in the pushing direction as "virtual anchors" to counteract the object's reaction forces. We refer readers to our video for the full manipulation.

Skiing. This example co-designs the anisotropic friction and control for a 3D robot with two skiboards on a slanted, uneven terrain. The objective for the robot is to execute speedy skiing across the terrain and end with a hard stop at a target. We co-optimize the anisotropic friction on the skiboards and the robot's controller to maintain balance. Randomly initialized friction before optimization can easily cause the robot to overshoot the target or lose balance, indicating that this task requires non-trivial co-optimization. The friction design and control policy learned from our pipeline exhibit substantial reorientation, with up to 103-degree changes in the friction to facilitate rapid skiing. The resulting locomotion demonstrates sharp turns followed by carefully planned braking (Fig. 1 and our video).

8.2 Examples in the Real World

Unit tests. We conduct unit tests on a three-link robot with two representative scale designs before reporting our real-world examples. We drive the robot with an open-loop square-wave control signal that alternates between $+90^\circ$ and -90° at 2 Hz, creating a motion pattern similar to a snake slithering. The first test presents the robot's locomotion with isotropic friction (Fig. 9): it struggles to move forward despite emulating a snake's slithering motion, a result also consistently reported by previous work [Qi et al. 2023]. This observation confirms the need to incorporate anisotropic friction for effective locomotion in this three-link robot.

The second test reports the robot's locomotion with all scales uniformly oriented towards the rear of the body, analogous to the organization of scales on a real snake [Hu et al. 2009]. Under the same open-loop control signal, the robot now gains an effective forward speed, emulating the slithering of a snake (Fig. 9). The observation confirms the efficacy of the robot hardware from our fabrication plan.

Turning. This example optimizes the anisotropic friction on a three-link robot to execute turning when driven by the same open-loop signal in the unit tests above. The robot demonstrates limited rotational capability with isotropic frictional scales (Fig. 9), as expected, because the open-loop signal is symmetric. We then optimize the anisotropic friction design using our pipeline by freezing the control and with an objective defined on the dot product between the robot's orientations at 0s and 10s. The optimized orientation of each link is $256^\circ, 138^\circ, 95^\circ$, respectively. They are significantly oriented in different directions, exhibiting a non-trivial layout. After deploying our 3D-printed scales with optimized anisotropic friction orientations on each link, the robot demonstrates strong rotational capability under the same control signal (Fig. 9), confirming the substantial influence of engineering anisotropic friction designs on robot locomotion.

We use a simulation experiment to study the impact of potential manufacturing errors. Taking the turning example as our benchmark, we randomly perturb scale orientations and simulate its performance (rotation angle) after the perturbation. The robot rotated 122° before perturbation and rotated 112° on average across 100 trials of up to 10° perturbation applied to each link. This implies that the robot's performance is robust to small perturbations in simulation.

Point tracking. Finally, we present an example that co-designs the anisotropic friction and closed-loop control for the three-link robot to track a sequence of points distributed along a straight line. We run our pipeline in Sec. 7 and assemble the 3D-printed scales with the optimized friction orientation on the robot links. We use a motion capture device (an Optitrack tri-lens camera) and send the robot's state to the trained neural network control policy running on a laptop in real time. The laptop then sends the control policy's output at 2 Hz to the robot.

Fig. 10 presents the locomotion of the robot after running our pipeline. It may be tempting to predict that the robot would execute a classic slithering motion, a straightforward solution to this point tracking task. However, the robot turns out to learn an unconventional strategy by twisting its body and moving laterally and successfully chases all targets in simulation. When transferred to the real world, the system is indeed able to accomplish point tracking. We attribute this strategy to the three different anisotropic friction orientations on the three links after optimization (196°, 293°, and 245°). The result highlights the potential of computational design to generate unconventional solutions with competitive performance compared with biomimetic designs inspired by natural evolution.

9 Conclusions, Limitations, and Future Work

Anisotropic friction is critical for animal locomotion, yet work on computationally designing anisotropic friction for robots is sparse. To our best knowledge, we present the first computational pipeline to address this challenging problem involving the interplay between anisotropic friction and control. Our evaluations have confirmed that our novel integration of trust-region and robot co-optimization is the key in our computational pipeline, leading to substantial performance improvement over one of the SOTA robot co-design methods. Furthermore, we have constructed LLM-driven robot co-design baselines and reported their highly competitive performance against SOTA methods, suggesting that LLMs warrant greater attention in future research on robot co-design. Examples in simulation have demonstrated that our pipeline can discover non-trivial anisotropic friction and control combinations that effectively accomplish various locomotion tasks. To validate our pipeline in real-world environments, we have collected, 3D-printed, and measured 55 scales with reliable anisotropic friction profiles on a custom-built hardware platform. Our real-world experiments have shown that our robots can turn, slither, and track points with optimized anisotropic friction and control guided by our pipeline, whereas robots with isotropic friction struggle to complete real-world locomotion, consistent with results reported in previous work [Qi et al. 2023].

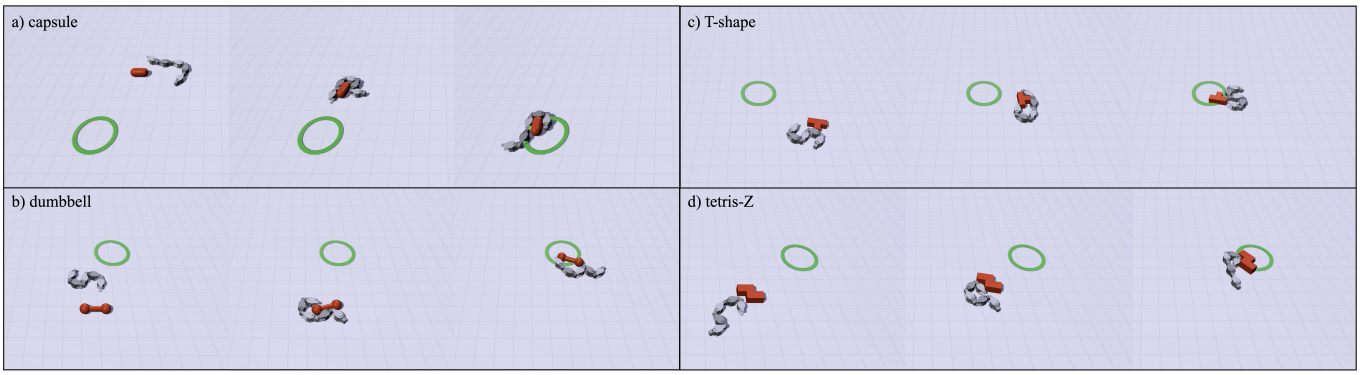

Fig. 8. Snapshots of the optimized robot manipulating four different objects: (a) *capsule*, (b) *dumbbell*, (c) *T-shape*, and (d) *tetris-Z*. The green circle indicates the target zone. The robot successfully adapts its undulation gait to the geometry of each object to complete the task.

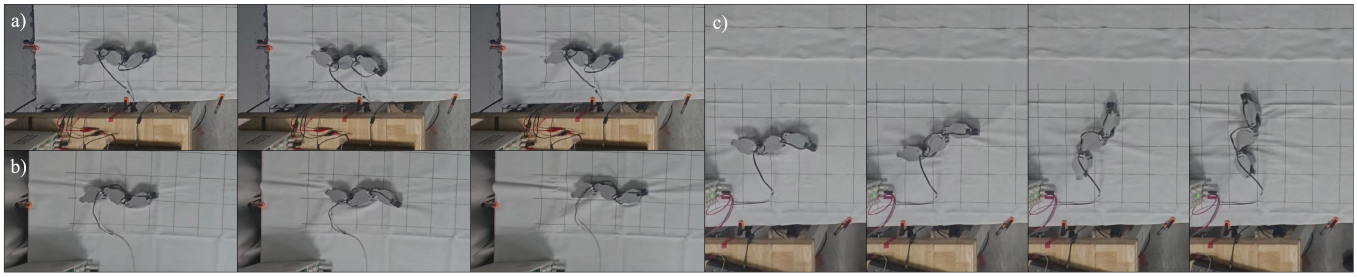

Fig. 9. Snapshots of our robot moving on a horizontal plane according to a slithering signal. a) Our robot struggles to move with isotropic scales, unable to move forward or turn. b) With scales naturally toward the rear of the body, our robot forwards with the slithering gait. c) After deploying scales optimized for the turning task, our robot exhibits directed turning.

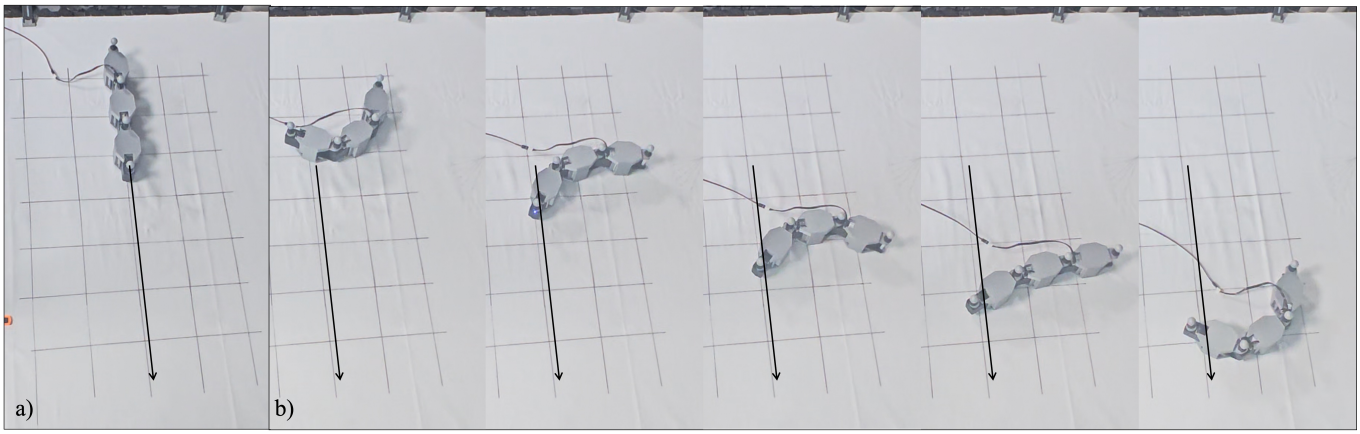

Fig. 10. *Point tracking*. our robot continuously tracks a target indicated by the black arrow on a horizontal plane using an unusual gait, after performing anisotropic friction and control optimization using our method. a) At the beginning, the robot is stationary and faces forward. b) Five snapshots during the motion, after the deployment of optimized scales and execution of optimized controller.

Our work has a few limitations. First, our current pipeline requires an anisotropic friction profile and focuses on optimizing its orientation on a given robot link-joint topology. Extending the pipeline to co-optimize robot topology and free-form friction profiles would possibly reveal more creative robot designs. Second, our usage of LLMs in the baselines is still preliminary. Given their competitive results against other SOTA methods, it is a promising direction to explore more advanced usage of LLMs in robot co-design. Third, our robot hardware consists of limited link and joint designs. Building robots with various link and joint types and optimizing them across

a wider range of tasks beyond locomotion would be a promising future direction.

Acknowledgments

We thank Bohui Wan and Jiamu Bu for their help with this paper. Tao Du acknowledges the support from Tsinghua University and the Shanghai Qi Zhi Institute Innovation Program.

References

Sheldon Andrews, Kenny Erleben, and Zachary Ferguson. 2022. Contact and friction simulation for computer graphics. In *ACM SIGGRAPH 2022 Courses*. 1–172.

Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. Openai gym. *arXiv preprint arXiv:1606.01540* (2016).

Tao Du, Adriana Schulz, Bo Zhu, Bernd Bickel, and Wojciech Matusik. 2016. Computational multicopter design. *ACM Trans. Graph.* 35, 6, Article 227 (Dec. 2016), 10 pages. https://doi.org/10.1145/2980179.2982427

David Eriksson, Michael Pearce, Jacob Gardner, Ryan D Turner, and Matthias Poloczek. 2019. Scalable global optimization via local Bayesian optimization. *Advances in neural information processing systems* 32 (2019).

Kenny Erleben, Miles Macklin, Sheldon Andrews, and Paul G Kry. 2020. The matchstick model for anisotropic friction cones. In *Computer Graphics Forum*, Vol. 39. Wiley Online Library, 450–461.

Carl Gans. 1962. Terrestrial locomotion without limbs. *American Zoologist* (1962), 167–182.

Agrim Gupta, Linxi Fan, Surya Ganguli, and Li Fei-Fei. 2022. Metamorph: learning universal controllers with transformers. In *International Conference on Learning Representations*. ICLR.

Agrim Gupta, Silvio Savarese, Surya Ganguli, and Li Fei-Fei. 2021. Embodied intelligence via learning and evolution. *Nature communications* 12, 1 (2021), 5721.

Yuanfeng Han, Hamidreza Marvi, and Metin Sitti. 2015. Fiberbot: A miniature crawling robot using a directional fibrillar pad. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 3122–3127.

David L. Hu, Jasmine Nirody, Terri Scott, and Michael J. Shelley. 2009. The mechanics of slithering locomotion. *Proceedings of the National Academy of Sciences* 106, 25 (2009), 10081–10085. https://doi.org/10.1073/pnas.0812533106 arXiv:https://www.pnas.org/doi/pdf/10.1073/pnas.0812533106

Jiaheng Hu, Julian Whitman, and Howie Choset. 2023. Glso: Grammar-guided latent space optimization for sample-efficient robot design automation. In *Conference on Robot Learning*. PMLR, 1321–1331.

Hyeonseong Jeon, Ainaz Eftekhar, Aaron Walsman, Kuo-Hao Zeng, Ali Farhadi, and Ranjay Krishna. 2025. Convergent Functions, Divergent Forms. *arXiv preprint arXiv:2505.21665* (2025).

Zhongying Ji, Shiyu Qin, Shuanhong Ma, Xin Jia, Xiaolong Wang, and Feng Zhou. 2022. Bio-inspired smart surface to achieve controllable locomotion through adjustable anisotropic friction. *Friction* 10, 8 (2022), 1180–1191.

Mina Konakovic Lukovic, Yunsheng Tian, and Wojciech Matusik. 2020. Diversity-guided multi-objective bayesian optimization with batch evaluations. *Advances in Neural Information Processing Systems* 33 (2020), 17708–17720.

Qiqin Le, Jiamu Bu, Yanke Qu, Bo Zhu, and Tao Du. 2024. Computational Biomimetics of Winged Seeds. *ACM Transactions on Graphics (TOG)* 43, 6 (2024), 1–13.

Minchen Li, Zachary Ferguson, Teseo Schneider, Timothy R Langlois, Denis Zorin, Daniele Panozzo, Chenfanfu Jiang, and Danny M Kaufman. 2020. Incremental potential contact: intersection-and inversion-free, large-deformation dynamics. *ACM Trans. Graph.* 39, 4 (2020), 49.

Haofei Lu, Zhe Wu, Junliang Xing, Jianshu Li, Ruoyu Li, Zhe Li, and Yuanchun Shi. 2025. BodyGen: Advancing Towards Efficient Embodiment Co-Design. In *The Thirteenth International Conference on Learning Representations*.

Pingchuan Ma, Tao Du, John Z. Zhang, Kui Wu, Andrew Spielberg, Robert K. Katzschmann, and Wojciech Matusik. 2021. DiffAqua: a differentiable computational design pipeline for soft underwater swimmers with shape interpolation. *ACM Trans. Graph.* 40, 4, Article 132 (July 2021), 14 pages. https://doi.org/10.1145/3450626.3459832

Hamidreza Marvi, Jacob Bridges, and David L Hu. 2013. Snakes mimic earthworms: propulsion using rectilinear travelling waves. *Journal of the Royal Society Interface* 10, 84 (2013), 20130188.

Hamidreza Marvi, Chaohui Gong, Nick Gravish, Henry Astley, Matthew Travers, Ross L Hatton, Joseph R Mendelson III, Howie Choset, David L Hu, and Daniel I Goldman. 2014. Sidewinding with minimal slip: Snake and robot ascent of sandy slopes. *Science* 346, 6206 (2014), 224–229.

Hamidreza Marvi and David L Hu. 2012. Friction enhancement in concertina locomotion of snakes. *Journal of the Royal Society Interface* 9, 76 (2012), 3067–3080.

Hamidreza Marvi, Gregory Meyers, Geoffrey Russell, and David L Hu. 2011. Scalybot: a snake-inspired robot with active control of friction. In *Dynamic Systems and Control Conference*, Vol. 54761. 443–450.

Jorge Nocedal and Stephen J Wright. 2006. *Numerical optimization*. Springer.

Xinda Qi, Tong Gao, and Xiaobo Tan. 2023. Bioinspired 3d-printed snakeskins enable effective serpentine locomotion of a soft robotic snake. *Soft Robotics* 10, 3 (2023), 568–579.

Ahmad Rafsanjani, Yuerou Zhang, Bangyuan Liu, Shmuel M Rubinstein, and Katia Bertoldi. 2018. Kirigami skins make a simple soft actuator crawl. *Science Robotics* 3, 15 (2018), eaar7555.

Miguel M Serrano, Alexander H Chang, Guangcong Zhang, and Patricio A Vela. 2015. Incorporating frictional anisotropy in the design of a robotic snake through the exploitation of scales. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 3729–3734.

Donghua Shen, Qi Zhang, Cunjin Wang, Xingsong Wang, and Mengqian Tian. 2021. Design and analysis of a snake-inspired crawling robot driven by alterable angle scales. *IEEE Robotics and Automation Letters* 6, 2 (2021), 3744–3751.

Junru Song, Yang Yang, Huan Xiao, Wei Peng, Wen Yao, and Feifei Wang. 2025. Laser: Towards diversified and generalizable robot design with large language models. In *The Thirteenth International Conference on Learning Representations*.

Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David Budden, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, et al. 2018. Deepmind control suite. *arXiv preprint arXiv:1801.00690* (2018).

Emanuel Todorov, Tom Erez, and Yuval Tassa. 2012. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 5026–5033. https://doi.org/10.1109/IROS.2012.6386109

Halvor T Tramsen, Stanislav N Gorb, Hao Zhang, Poramate Manoonpong, Zhendong Dai, and Lars Heepe. 2018. Inversion of friction anisotropy in a bio-inspired asymmetrically structured surface. *Journal of The Royal Society Interface* 15, 138 (2018), 20170629.

Nobuyuki Umetani, Yuki Koyama, Ryan Schmidt, and Takeo Igarashi. 2014. Pteromys: Interactive design and optimization of free-formed free-flight model airplanes. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–10.

Tsun-Hsuan Wang, Pingchuan Ma, Andrew Everett Spielberg, Zhou Xian, Hao Zhang, Joshua B. Tenenbaum, Daniela Rus, and Chuang Gan. 2023. SoftZoo: A Soft Robot Co-design Benchmark For Locomotion In Diverse Environments. In *The Eleventh International Conference on Learning Representations*. https://openreview.net/forum?id=Xyme9p1rpZw

Ye Yuan, Yuda Song, Zhengyi Luo, Wen Sun, and Kris Kitani. 2022. Transform2Act: Learning a Transform-and-Control Policy for Efficient Agent Design. In *International Conference on Learning Representations*.

Allan Zhao, Jie Xu, Mina Konaković-Luković, Josephine Hughes, Andrew Spielberg, Daniela Rus, and Wojciech Matusik. 2020. Robogrammar: graph grammar for terrain-optimized robot design. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–16.