

Swimming Under Constraints: A Safe Reinforcement Learning Framework for Quadrupedal Bio-Inspired Propulsion

Xinyu Cui^{1,2,3*}, Fei Han^{1,4*†}, Hang Xu¹, Yongcheng Zeng^{2,3}, Luoyang Sun^{2,3}, Ruizhi Zhang², Jian Zhao⁵, Haifeng Zhang², Weikun Li¹, Hao Chen¹, Jun Wang⁶, Dixia Fan^{1†}

Abstract—Bio-inspired aquatic propulsion offers high thrust and maneuverability but is prone to destabilizing forces such as lift fluctuations, which are further amplified by six-degree-of-freedom (6-DoF) fluid coupling. We formulate quadrupedal swimming as a constrained optimization problem that maximizes forward thrust while minimizing destabilizing fluctuations. Our proposed framework, Accelerated Constrained Proximal Policy Optimization with a PID-regulated Lagrange multiplier (ACPPO-PID), enforces constraints with a PID-regulated Lagrange multiplier, accelerates learning via conditional asymmetric clipping, and stabilizes updates through cycle-wise geometric aggregation. Initialized with imitation learning and refined through on-hardware towing-tank experiments, ACPPO-PID produces control policies that transfer effectively to quadrupedal free-swimming trials. Results demonstrate improved thrust efficiency, reduced destabilizing forces, and faster convergence compared with state-of-the-art baselines, underscoring the importance of constraint-aware safe RL for robust and generalizable bio-inspired locomotion in complex fluid environments.

Index Terms—Underwater Bio-inspired Robots; Safe Reinforcement Learning; Motion Planning; Force Sensing

I. INTRODUCTION

Bio-inspired aquatic propulsion systems, such as flapping foils, undulating fins, and webbed paddles have attracted significant attention for their ability to generate large instantaneous thrust [1], [2], achieve high propulsive efficiency, and enable agile maneuverability [3], [4], [5]. However, unlike conventional screw or jet drives, the asymmetric morphology and unsteady interactions of these propulsors often introduce destabilizing effects, including lift fluctuations, pitch oscillations, and lateral forces at peak thrust [6], [7]. Such disturbances elevate energy consumption, compromise stability, and degrade mission performance.

* Equal contribution, the order is determined by flipping a coin.

† Corresponding to: {hanfei, fandixia}@westlake.edu.cn

This work is supported by the Key Research and Development Program of Zhejiang Province (Grant No. 2023C03133), the Young Scientists Fund of the National Natural Science Foundation of China (Grant No. 52401393), the Natural Science Foundation of Hangzhou, Zhejiang Province, China (Grant No. 2025SZRJ1217), and the Hangzhou Postdoctoral Daily Expenses Funding (Grant No. 103120046582502).

¹ FSI lab, Westlake University, Hangzhou, 310030, China.

² Institute of Automation, Chinese Academy of Sciences, Beijing, 100190, China.

³ School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing 100049, China

⁴ Zhejiang University, Hangzhou, 310027, China.

⁵ Zhongguancun Academy, Beijing, 100094, China

⁶ Department of Computer Science, University College London, London WC1E 6BT, United Kingdom

This challenge extends to all fluid-coupled vehicles operating underwater, on surface, or in air. Large pitch and roll excursions, with disturbances induced by wave or vortex, increase drag and impair maneuverability [8], [9], [10]. Unsteady fluid–structure interactions also generate destabilizing moments that reduce efficiency and control authority [11]. Therefore, locomotion design needs to constrain unwanted oscillatory forces while maximizing thrust.

While pre-defined gaits with fixed parameters can suppress some of these effects, they under-utilize the nonlinear dynamics of unsteady hydrodynamics and fail to fully leverage the robot’s embodiment [12], [13], [14]. Reinforcement learning (RL) provides a promising alternative by enabling propulsors to autonomously discover control policies that exploit embodiment and hydrodynamics. However, naive RL exploration typically results in instability and inefficient convergence when balancing the dual objectives of maximizing thrust while minimizing unwanted forces [15]. Safe RL offers a principled framework to address this challenge: explicit constraints guide policy updates to ensure high performance without violating safety or stability bounds [16]. While safe RL has seen success in ground and aerial robotics, its application to aquatic learning remains largely unexplored, where strong fluid coupling and costly experimentation make stability guarantees especially critical.

This paper uses fast forward swimming in a bio-inspired quadrupedal robot to validate our proposed safe reinforcement learning algorithm. We formulate gait learning as a constrained optimization problem that maximizes thrust while bounding lift oscillations—the primary destabilizing factor. To solve this, we propose Accelerated Constrained PPO with a PID-regulated Lagrange multiplier (ACPPO-PID). Our approach enhances on-hardware learning by dynamically enlarging the PPO [17] clip range under constraint satisfaction, and employing cycle-wise geometric aggregation for stability. The policy is seeded via Imitation Learning, refined through real-time towing-tank reinforcement learning, and deployed on a free-swimming robot using diagonal-phase paddling. Experiments show ACPPO-PID significantly outperforms state-of-the-art baselines in thrust efficiency, lift suppression, and convergence speed.

Contributions. This paper makes three main contributions:

- We formulate quadrupedal swimming as a constrained thrust optimization problem, decoupling full-robot op-

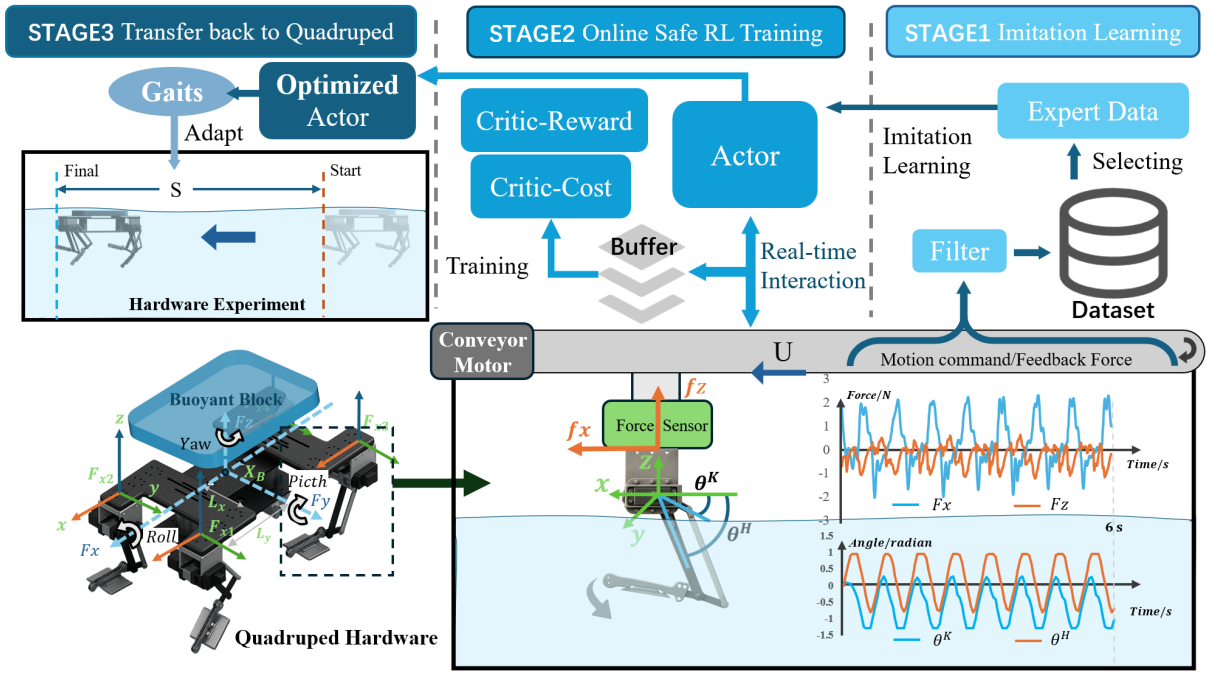


Fig. 1. Proposed framework: imitation learning initializes a periodic gait from predefined motions while sensor feedback is Kalman-filtered; safe RL fine-tuning accelerates on-hardware convergence under stability constraints; and the resulting one-cycle paddle is transferred to diagonal limb pairs with a half-cycle phase offset to enable smooth and stable free-swimming.

timization into single-limb gait optimization.

- We propose ACPPO-PID, a safe RL framework that balances broad exploration with strict constraint enforcement, accelerating on-hardware learning.
- We validate the approach through real-world towing-tank and quadruped swimming experiments, demonstrating improved thrust, stability, and convergence over state-of-the-art baselines.

II. RELATED WORK

A. RL for Bio-inspired Robotic Locomotion

RL has emerged as a powerful tool for optimizing locomotion in bio-inspired robots. In terrestrial settings, RL has been successfully applied to legged robots for robust walking and running [18], [19], often learning directly from proprioceptive sensing and reward signals related to forward progress and energy efficiency. In aquatic environments, RL has been employed to optimize swimming gaits in robotic fish [1], [20], undulating fins [5], and other bio-inspired propulsors [9].

However, these methods typically focus on maximizing thrust or speed without explicitly accounting for stability constraints or unwanted forces like lift in a forward-moving task. Additionally, their reliance on simulation may not accurately capture complex fluid-structure interactions, creating sim-to-real gaps. In contrast, our work emphasizes on-hardware learning with explicit stability constraints, directly addressing the reality gap and improving policy practicality.

B. Safe RL for Constrained Optimization

Safe RL addresses constrained optimization through several paradigms. Local policy search methods like CPO

[21] extend TRPO [22] with second-order approximations, while FOCOPS [23] reduces computational overhead through feasible-trust region intersections. Projection-based approaches such as PCPO [24] and CUP [25] optimize within trust regions before projecting policies into feasible spaces, often yielding conservative outcomes.

Primal-dual methods transform constraints into unconstrained Lagrangian forms, with RCPO [26] adjusting multipliers for stability and CPPO-PID [27] incorporating PID controllers [28] for improved convergence. While primal-dual approaches offer computational efficiency suitable for on-hardware learning, they exhibit performance instability and require careful hyperparameter tuning. Our method addresses these limitations by enhancing exploration when safe while improving training stability through cycle-wise aggregation. This approach can also steer policies toward constraint satisfaction when violations persist, reducing hyperparameter sensitivity for practical deployment.

III. PRELIMINARIES

To validate the optimization framework and coordination strategy, we developed two complementary systems: a **single-leg towing platform** for learning physical constraints in controlled flows, and a **quadrupedal underwater robot** for validating gait generalization in free-swimming scenarios. This setup forms a closed loop from isolated actuator learning to embodied whole-robot validation.

A. Quadrupedal Underwater Robot and System Dynamics

As illustrated in the schematic diagram of the robotic system (Fig. 1, bottom left), the quadrupedal underwater

robot comprises a support plate, four independent single-leg actuators, and modular buoyancy units, with overall dimensions of 480 mm × 250 mm × 88 mm.

The dynamic model of the robot is derived based on its three-dimensional geometry, incorporating both the thrust forces and intrinsic joint torques generated by each leg actuator. Within this model, X_B denotes the center of buoyancy. The position of each actuator is defined in a body-fixed coordinate system with origin at X_B , represented as (L_x, L_y, h) , where L_x and L_y are the horizontal offsets from X_B , and h indicates the vertical eccentricity relative to the central rotational axis. For gait synthesis and analysis, the four foot actuators are abstracted as two diagonal pairs: $\{P_1, P_4\}$ and $\{P_2, P_3\}$, simplifying the control strategy while preserving dynamic characteristics.

Diagonal coordination and constrained optimization: Owing to the approximate geometric symmetry of the four limbs about both longitudinal and lateral axes, a diagonally symmetric coordination strategy is adopted. The four limbs are grouped into two diagonal pairs, and each pair is driven either in phase or with a half-cycle offset depending on the gait design. This strategy provides two benefits: (i) in-plane force components tend to cancel out in yaw due to symmetry, and (ii) the overlap of thrust peaks and troughs between diagonals smooths the net force profile, thereby reducing oscillations in propulsion and suppressing attitude disturbances. Since the limb kinematics are mechanically confined to the sagittal plane, lateral hydrodynamic forces are inherently minimized, with any residuals effectively neutralized by the symmetric coordination. As the limbs generate the primary propulsion and their motion patterns are assumed similar, the gait optimization problem of the quadruped can thus be reduced to optimizing a single representative limb, which is subsequently deployed to the whole robot via diagonal coordination.

In the validation experiments, the same single-leg motion strategy is mapped to diagonal pairs $\{P_1, P_4\}$ and $\{P_2, P_3\}$. Due to diagonal symmetry, the legs within each pair maintain synchronized motion and generate identical forces and moments. Let \mathbf{F}_1 and $\boldsymbol{\tau}_1$ represent the force and moment contribution from a single leg in the first pair, and $\mathbf{F}_2, \boldsymbol{\tau}_2$ for the second pair. With a vertical offset h relative to the center of buoyancy, the net force and moments are derived as $\mathbf{F}_{\text{total}} = 2(\mathbf{F}_1 + \mathbf{F}_2) = (f_X, f_Y, f_Z)$, $M_X = 2(\tau_{1x} + \tau_{2x}) - hf_Y$, $M_Y = 2(\tau_{1y} + \tau_{2y}) + hf_X$, and $M_Z = 2(\tau_{1z} + \tau_{2z})$. This formulation demonstrates that yaw moments from in-plane forces cancel under diagonal symmetry, while the pitch and roll moments are directly coupled with the total horizontal thrust and the offset h .

B. Two-DOF Single-Limb Platform and Learning Objective

As shown in Fig. 1(bottom right), a single-limb hardware platform with two degrees of freedom (DoF) is designed for controlled evaluation. The limb consists of a hip flexion/extension (HFE) joint and a knee flexion/extension (KFE) joint, mechanically coupled through a four-bar linkage. This configuration yields a compact structure while allowing

large-amplitude, coordinated swinging motions. Each joint is independently actuated by a servo motor, and the motion is transmitted via rigid links to a web. A six-axis force/torque sensor is mounted near the web tip to measure instantaneous three-dimensional force and torque components, with emphasis on forward thrust F_x , vertical lift F_z , and pitching moment M_y .

Experimental setup: The platform is tested in a towing tank, where the limb is mounted on a horizontal carriage translated at constant speed and fixed immersion depth while executing prescribed joint trajectories. Given commanded joint angles θ^H and θ^K , the four-bar mechanism drives the web through a full oscillation cycle. During each cycle, the force/torque sensor synchronously records the time histories of F_x , F_z , and the associated moment components. This setup provides a repeatable and high-resolution means of evaluating hydrodynamic performance under different motion policies.

Learning objective: Based on the force-motion characteristics of the symmetrically arranged four legs of the quadruped robot platform, during diagonal-symmetric motion, the yaw moment M_Z and roll moment M_X are essentially zero, while the pitch moment remains consistently positive. As a result, in an unconstrained scenario, the robot moves in a straight line with a 'pitch-up' posture. Its motion effect is mainly influenced by the forward thrust f_X and the vertical force f_Z , which causes vertical oscillations. The control objective is formulated as maximizing the time-averaged forward thrust F_x while simultaneously suppressing both the magnitude and variability of the lift F_z . This design explicitly reduces ineffective oscillatory work and improves locomotor stability. Crucially, the optimized single-limb policy can be directly transferred to the quadrupedal robot via diagonal coordination, thereby achieving high propulsion efficiency together with robust whole-body stability in free-swimming conditions.

C. Problem Formulation

Subsequently, we consider the optimization of a flapping limb to produce higher thrust while maintaining low oscillations in lift forces. To fully explore the hydrodynamic, we adopt a more agile end-to-end scheme instead of representing the limb's motion in a parametric sinusoidal form. We then formulate a reward function proportional to thrust while treating oscillation as a constraint. Formally, we cast control of a single limb as a Constrained Markov Decision Process (CMDP) [29] $\langle \mathcal{S}, \mathcal{A}, P, R, C \rangle$, where:

- State $s_t \in \mathcal{S}$: measured quantities of the limb, including joint phase, angular velocity, and the measured values of the sensors.
- Action $a_t \in \mathcal{A}$: the angle change of the joints in one control step.
- Transition $P(s_{t+1} | s_t, a_t)$: unknown dynamics, learned from interaction.
- Reward $R(s_t, a_t, s_{t+1})$: proportional to forward thrust component, denoted r_t .

- Cost $C(s_t, a_t, s_{t+1})$: lift non-cancellation across a half motion cycle, denoted c_t . With a cycle of H steps, we define the instantaneous cost as $c_t = |F_{z,t} + F_{z,t-H/2}|$.

Let $\tau = (s_0, a_0, s_1, a_1, \dots)$ be a trajectory generated by policy $\pi(a | s)$ and discount factor $\gamma \in (0, 1)$. The expected discounted return and cost are

$$\begin{aligned} J(\pi) &= \mathbb{E}_{\tau \sim \pi} \left[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t, s_{t+1}) \right], \\ J_C(\pi) &= \mathbb{E}_{\tau \sim \pi} \left[\sum_{t=0}^{\infty} \gamma^t C(s_t, a_t, s_{t+1}) \right]. \end{aligned} \quad (1)$$

The learning objective is

$$\pi^* \in \arg \max_{\pi} J(\pi) \quad \text{s.t.} \quad J_C(\pi) \leq d, \quad (2)$$

where $d > 0$ bounds the allowable lift fluctuation.

To solve this constrained optimization problem, we introduce the Lagrangian function, defined as

$$\begin{aligned} \mathcal{L}(\pi, \lambda) &= J(\pi) + \lambda(d - J_C(\pi)), \\ &= \mathbb{E}_{\tau \sim \pi} \left[\sum_{t=0}^{\infty} \gamma^t (R(s_t, a_t, s_{t+1}) - \lambda C(s_t, a_t, s_{t+1})) \right] + \lambda d, \end{aligned} \quad (3)$$

where $\lambda \geq 0$ is the Lagrange multiplier. This formulation transforms the constrained problem into an equivalent unconstrained saddle-point optimization problem: $\min_{\lambda \geq 0} \max_{\pi} \mathcal{L}(\pi, \lambda)$.

IV. METHODOLOGY

We propose a three-stage safe policy optimization framework, as illustrated in Fig. 1. In the first stage, a periodic paddle motion is initialized through demonstration-driven imitation learning, providing a stable starting point for subsequent training. In the second stage, the policy is refined via on-hardware safe reinforcement learning in a towing tank, enabling constraint-aware adaptation under real hydrodynamic interactions. In the final stage, the optimized policy is transferred to the quadrupedal robot by recording a single cycle of joint trajectories and deploying it with diagonal-phase coordination, thereby achieving stable and efficient free-swimming in validation trials.

A. Imitation Learning

To reduce the wall-clock time required for on-hardware training, we first initialize the policy using **Imitation Learning (IL)** on curated demonstrations. These demonstrations are generated by parameterizing limb trajectories as sinusoids and performing a brute-force (BF) search over frequency, amplitude, and phase. The prescribed leg motion is expressed as:

$$\begin{cases} \theta^H(t) = A_{\theta H} \sin(2\pi ft) + \theta_0^H, \\ \theta^K(t) = A_{\theta K} \sin(2\pi ft + \phi) + \theta_0^K, \end{cases} \quad (4)$$

where, $A_{\theta H}$ and $A_{\theta K}$ represent the angles amplitude of the HFE joint and the KFE joint. The motion frequency is denoted by f , and ϕ is the phase difference between the HFE and KFE joints. θ_0^H and θ_0^K represent the initial offset in motion.

TABLE I
EXPERIMENTAL PARAMETER RANGES

Parameter	$A_{\theta H}$ (rad)	$A_{\theta K}$ (rad)	f (Hz)	ϕ	θ_0^H / θ_0^K
Range	$[\pi/6, \pi/3]$	$[\pi/12, \pi/4]$	$[0.3, 0.6]$	$[0, \pi]$	$[\pi/4, 5\pi/4]$

To generate expert demonstrations, we employ **Latin Hypercube Sampling (LHS)** over the parameter space in Eq. 4 and Table I, yielding over five thousand sinusoidal trajectories ranked by time-averaged thrust, with the lowest-lift subset retained as the demonstration set. Additionally, the best-performing sinusoidal trajectory from the search is extracted and denoted as **BF**, which serves as the parameterized baseline in our experiments. Demonstration rollouts are stored as state-action pairs and are used to pretrain the policy via IL, enabling rapid convergence to a stable periodic gait. This provides a safe, efficient initialization for subsequent on-hardware safe RL finetuning.

For both the policy and value functions, we adopt a **Transformer** architecture. Unlike recurrent models that compress history into a fixed hidden state which potentially losing high-frequency fluid details. Transformer's self-attention mechanism enables direct access to the entire observation window. This design is critical for aquatic propulsion, where hydrodynamic forces exhibit significant delays that recurrent bottlenecks often fail to capture. A sequence of recent observations is processed through the encoder, and the final embedding is decoded by a multilayer perceptron (MLP) to generate continuous actions, providing a robust foundation for stability-aware control.

B. Safe RL with accelerated safe exploration

We estimate the principal paddle frequency from \mathcal{F}_z , a sequence of F_z observed in the replay buffer via a discrete Fourier transform (DFT). Let the control rate be $f_s = 20$ Hz. We remove low-frequency drift below 0.1 Hz and restrict the search to $[0.1, 5]$ Hz to reflect reciprocating paddles. The dominant frequency is

$$f^* = \arg \max_{f \in [0.1, 5]} |\text{DFT}(\mathcal{F}_z)|, \quad (5)$$

and the cycle (in steps) is $H = \left\lfloor \frac{f_s}{f^*} \right\rfloor$. To promote lift cancellation under the half-cycle phase offset between limb groups, we have the instantaneous cost as $c_t = |F_{z,t} + F_{z,t-H/2}|$, which directly penalizes residual lift after superposing signals separated by half a cycle.

To solve the constrained optimization problem introduced in the problem formulation, we employ an iterative algorithm that alternates between policy optimization and Lagrange multiplier updates. The process proceeds over iterations indexed by $k = 0, 1, 2, \dots$, starting with an initial $\lambda_0 \geq 0$ (typically set to 0).

At each iteration k , we first optimize the policy parameters θ to maximize the expected Lagrangian return using the current multiplier λ_k :

$$\max_{\theta} \mathbb{E} \left[\sum_t \gamma^t (r_t - \lambda_k c_t) \right]. \quad (6)$$

Next, we collect trajectories under the updated policy π_θ , compute an empirical estimate of the cost return $\hat{J}_C^{(k)}$, and measure the constraint violation $g_k = \hat{J}_C^{(k)} - d$. Finally, we update the multiplier for the next iteration using PID control as defined in CPPO-PID [27]:

$$\lambda_{k+1} = \left[\lambda_k + K_P g_k + K_I \sum_{i=0}^k g_i + K_D (g_k - g_{k-1}) \right]_+,$$

where $[\cdot]_+$ projects onto $[0, \infty)$, and K_P, K_I, K_D are hyperparameters for proportional, integral, and derivative control, respectively. This PID mechanism accelerates convergence to feasible policies by responsively adapting λ to observed constraint violations. We maintain separate value functions for reward and cost to compute Generalized Advantage Estimates A_t^r and A_t^c , and form a Lagrangian advantage $A_t^\lambda = \bar{A}_t^r - \lambda \bar{A}_t^c$ with normalized advantages \bar{A}_t^r and \bar{A}_t^c .

Standard CPPO-PID uses a symmetric clip $[1 - \epsilon, 1 + \epsilon]$ to ensure conservative updates, which can slow on-hardware learning. We propose Accelerated CPPO-PID by integrating several advances. Inspired by [30], we conditionally enlarge only the upper clip bound when the estimated A_t^r is positive and A_t^c is non-positive to enhance wider exploration without sacrificing safety. Formally, with $\rho_t = \frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{\text{old}}}(a_t|s_t)}$, we define an asymmetric upper bound:

$$\epsilon_t^+ = \begin{cases} \epsilon_{\text{hi}} = 0.28, & \text{if } A_t^r > 0 \wedge A_t^c \leq 0 \wedge \text{ep} \geq \text{ep}_{\text{warm}}, \\ \epsilon = 0.2, & \text{otherwise,} \end{cases} \quad (7)$$

where ep is the training episode, ep_{warm} is a brief warm-up after which value net fits stabilize. The step-wise actor surrogate becomes:

$$L_{\text{step}}(\theta) = -\mathbb{E} \left[\min \left(\rho_t A_t^\lambda, \text{clip}(\rho_t, 1 - \epsilon, 1 + \epsilon_t^+) A_t^\lambda \right) \right]. \quad (8)$$

Locomotion episodes comprise repeated cycles. To align updates with cycle-level performance, we aggregate importance ratios over each detected cycle p of length H . We operate in the log domain by defining a logged signed importance ratio $\iota_t = \log \rho_t \cdot \text{sign}(A_t^\lambda)$. Inspired by recent work on geometric aggregation [31], we compute a clipped geometric mean for the cycle-wise importance ratio:

$$\tilde{\rho}_p = \exp \left[\frac{1}{|H|} \sum_{t \in p} \left(\min(\iota_t, \text{clip}(\iota_t, -\epsilon_p, \epsilon_p)) \cdot \text{sign}(A_t^\lambda) \right) \right]. \quad (9)$$

The geometric mean is less sensitive to outliers and provides a smoother cycle-level importance metric. To encourage exploration at this cycle level, we adopt a larger clipping value ϵ_p of up to 0.4. The corresponding cycle-level surrogate is:

$$L_{\text{cyc}}(\theta) = -\mathbb{E} \left[\tilde{\rho}_p A_t^\lambda \right], \text{ where } t \in p. \quad (10)$$

The final actor objective blends local and global views,

$$L_{\text{actor}}(\theta) = \alpha L_{\text{step}}(\theta) + (1 - \alpha) L_{\text{cyc}}(\theta), \quad (11)$$

where $\alpha \in (0, 1)$ balances step-wise stability and cycle-level alignment. This hybrid objective provides a multi-faceted

view that enhances both policy exploration and training stability.

Moreover, letting \bar{A}_p denote the mean advantage over a cycle $\frac{1}{|H|} \sum_{t \in p} A_t^\lambda$, we have the actor gradient for the cycle-wise loss.

$$\nabla_\theta L_{\text{cyc}}(\theta) = -\mathbb{E}_p \left[\bar{A}_p \tilde{\rho}_p \frac{1}{|H|} \sum_{t \in p} \mathbf{1}\{\iota_t \leq \epsilon_p\} \chi_t \right]. \quad (12)$$

The geometric mean aggregation not only captures the policy change over the entire cycle but is robust to outliers, filtering spikes smoothly to provide a stable, global importance metric. Crucially, the gradient aggregation combines local gradients weighted by this stable, cycle-level statistic. Letting \mathbf{g} denote the average gradient $\sum_{t \in p} \chi_t$, the projection of this gradient onto the average policy gradient direction within a cycle is $\nabla_\theta L_{\text{cyc}}(\theta) \cdot \mathbf{g} = \bar{A}_p \tilde{\rho}_p \|\mathbf{g}\|^2$. Because $\tilde{\rho}_p > 0$ and $\|\mathbf{g}\|^2 \geq 0$, the sign of the dot product between the average gradient and the aggregated one matches the sign of \bar{A}_p .

This approach is effective because \bar{A}_p serves as a reliable performance metric. Since the goal is to optimize sustained performance over an infinite horizon and each step's thrust contribution is approximately uniform, we expect the discount factor γ to approach 1. In this setting, the value function reflects the long-term expected return. The advantage A_t at each step measures the immediate benefit of an action relative to this long-term average. By averaging these step-wise advantages over a full locomotion cycle to compute \bar{A}_p , we create a metric that filters out noisy, single-step evaluations. This cycle-averaged advantage robustly indicates whether the sequence of actions within that cycle contributed positively or negatively to the overall objective, providing a stable learning signal.

This stability provides a significant advantage for practical safe reinforcement learning. It reduces the reliance on extensive hyperparameter tuning, which is a common requirement for other methods but is often impractical for on-hardware experiments. In our method, the Lagrange multiplier λ increases when constraints are violated, biasing A_t^λ negatively to prioritize cost reduction. When the policy persistently violates safety constraints, our cycle-level objective provides global guidance towards safer regions, avoiding the erratic updates typical of step-wise approaches. This mitigates the tendency of CPPO-PID to temporarily sacrifice constraint satisfaction for reward, thereby reducing hyperparameter sensitivity and accelerating safe exploration.

C. Policy Transferring to Quadruped

After online safe RL training, the learned policy is transferred to the original quadrupedal swimming robot. Because the gait is periodic, the policy is first executed in the tow tank in inference mode, and one full cycle of the commanded joint trajectories is recorded as a gait primitive. This cycle is then applied to all four limbs as a baseline description of the paddle. Power and communications are provided via external cables for centralized control and data acquisition.

To preserve symmetry and reduce reaction lift fluctuation, the four limbs are partitioned into two diagonal pairs. The

two pairs execute the same recorded cycle with a phase difference of half a cycle. This alignment allows peaks and troughs of pairwise forces to overlap in time, which smooths the aggregate force profile and suppresses oscillations.

The resulting gaits from different trained policies are evaluated on the quadrupedal robot. Each gait is tested for the farthest forward distance traveled within a fixed time window, which quantifies propulsion efficiency, oscillation suppression, and stability. Moreover, supplementary videos are provided for clearer and more intuitive visualization of the swimming behavior.

V. EXPERIMENTS AND RESULTS

To evaluate the learning capability of the proposed framework and its suitability for real-world deployment, we conduct two complementary sets of experiments. First, gait optimization is performed in a towing tank, enabling controlled assessment of single-limb dynamics and constraint-aware learning. Subsequently, the optimized gaits are transferred to the quadrupedal robot and validated in free-swimming trials, demonstrating their effectiveness in achieving stable and efficient locomotion.

A. Gait Optimizing

We train the safe RL policy using towing data and on-policy rollouts. Validation trials in the towing tank assess thrust maintenance and lift reduction. The towing carriage advances at a constant speed of 0.15 m/s. During towing, sensor signals are Kalman-filtered and downsampled to 20 Hz to achieve frequency unification with the limb controller. The limb has two DoFs; for each DoF, the maximum swing amplitude is 20° . For each training episode, we collect about 360 control steps and hyperparameters are kept consistent across experiments. The ACPPO-PID parameter α is set to 0.2 and ep_{warm} is set to 10.

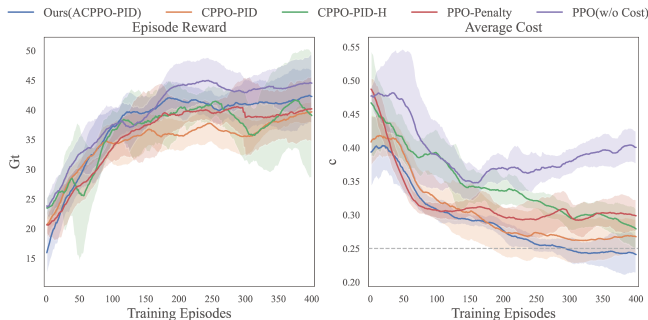


Fig. 2. The training curves (reward and average cost) of the baselines over 400 episodes, each method is trained in three random seeds. The gray line marks the cost limit.

1) *Baselines*: We compare Our method(ACPPO-PID) against the following to demonstrate benefits in both constraint satisfaction and learning efficiency:

- CPPO-PID: a SOTA safe RL algorithm,
- CPPO-PID-H: a naive high-clip variant of CPPO-PID using a fixed clip $[1-\epsilon, 1+\epsilon_{hi}]$ to encourage exploration irrespective of safety.

- PPO-Penalty: instead of optimizing a constrained objective, incorporates cost as a scalar penalty on reward $r \leftarrow r - 0.5 \cdot c$.
- PPO without cost: ignores the cost signal and trains a standard PPO policy without constraints.

All methods are trained from a policy initialized by IL.




Fig. 3. Ablations removing the cycle-level objective, conditional high clipping, or imitation learning quantify each component’s contribution to stability and data efficiency. Each ablation is trained in three random seeds. The gray line marks the cost limit.

2) *Ablation Study*: Our method accelerates learning in three respects. To assess their contributions, we conducted ablations by removing each component in turn and compare the resulting variants:

- without cycle loss: removes the cycle-aligned objective L_{cyc} , training solely with the step-wise surrogate.
- without asymmetric clipping: enforces symmetric clipping $[1-\epsilon, 1+\epsilon]$ at all steps, disabling the conditional upper-bound enlargement used for safe, wider exploration.
- without IL: omits the imitation learning initialization and trains the policy from scratch on-hardware.

3) *Metrics*: We evaluate all methods by undiscounted episode reward and average cost (Table II). For each method, we train three independent models with different random seeds. After 400 training episodes, each model is evaluated in inference mode over three rollouts; we report the mean and standard deviation across nine runs per method. The policies trained from scratch fail to converge to a feasible motion and is excluded from the final comparison. We also include a sinusoidal paddling policy with the best parameters found by BF (denoted BF), evaluated over three rollouts.

TABLE II

SINGLE LIMB EVALUATED IN A TOW TANK

Algorithm	Episode Reward	Average Cost
CPPO-PID	39.58 ± 3.93	0.249 ± 0.021
CPPO-PID-H	37.10 ± 10.84	0.261 ± 0.023
PPO-Penalty	43.46 ± 3.91	0.314 ± 0.021
PPO (w/o Cost)	48.58 ± 5.15	0.391 ± 0.039
Ours (ACPPO-PID)	45.19 ± 3.83	0.235 ± 0.019
Ours (w/o Cycle Loss)	43.35 ± 5.92	0.271 ± 0.027
Ours (w/o Asym Clip)	38.66 ± 5.56	0.212 ± 0.026
BF	26.31 ± 1.02	0.332 ± 0.008

4) *Advantages of ACPPO-PID: Accelerated Learning*. Fig 2 presents the learning curves of episodic reward and

average cost. Table II summarizes the performance of three random seeds evaluated three times each. PPO without cost attains the highest reward but incurs severe constraint violations. Among cost-aware algorithms, ACPPO-PID achieves the best reward-cost trade-off within 400 training episodes. As an accelerated and stable advance of CPPO-PID, it attains the highest reward while maintaining the lower cost.

Comparing with a naive variant of encouraging exploration by simply widening the clipping range to $[1 - \epsilon, 1 + \epsilon_{hi}]$, denoted as CPPO-PID-H, we found it suffered from training instability. As the aggregated advantage A_t^λ is computed from the normalized reward and cost advantages, it only ranks step-wise relative advantage within a training batch. In contrast, our stepwise selection applies stricter filtering so that only genuinely advantageous updates proceed, and the cycle-wise smoothing over an entire paddle further promotes safe exploration. Moreover, as shown in Fig. 3, it is evident that imitation learning is essential. Without an initial policy, 400 on-hardware episodes are insufficient for the limb to discover a feasible paddling motion.

Stability and Constraint Satisfaction. As shown in Fig. 2 and Fig. 3, the proposed method achieves consistently higher stability, whereas other approaches with larger clipping ranges exhibit pronounced fluctuations. Moreover, our method effectively suppresses cost while sustaining high performance. This advantage stems from the dual-component design of ACPPO-PID. When safety constraints are repeatedly violated, the cycle-level objective steers the entire policy update toward constraint satisfaction. In parallel, the step-wise mechanism more strictly identifies safe and advantageous actions, applying a larger update ratio to reinforce them. The synergy of these two mechanisms enables safe yet accelerated exploration, thereby validating the effectiveness of the proposed improvements.

B. Quadruped Swimming

For quantitative validation, free-swimming experiments are conducted in a 4.0 m \times 1.5 m water tank. The robot is tested under various gaits, each running for 6 s and repeated three times. Swimming distance serves as the efficiency metric: a longer average distance indicates superior time-averaged thrust and reduced ineffective work during the swing phase, thereby indirectly validating the proposed single-leg optimization method.




Fig. 4. Forward swimming performance under two representative gaits per algorithm, averaged over three trials. ACPPO-PID achieves the highest scores with lower variance, while CPPO-PID, PPO, and BF show reduced thrust and greater sensitivity to gait differences.

As shown in Fig. 4, we deploy two representative gaits per algorithm on the quadruped robot, each executed three times and averaged. The proposed **ACPPO-PID** achieves the best performance, outperforming the constrained baseline *CPPO-PID* by about 7%, the unconstrained *PPO* by 8–9%, and the brute-force search (*BF*) by over 19–20%. These results confirm that ACPPO-PID not only produces stronger forward thrust but also maintains more consistent performance across gaits, achieving a better balance of propulsion efficiency and stability.




Fig. 5. Comparison of F_x^{mean} , F_z^{mean} , and F_z^{var} across algorithms, showing ACPPO-PID achieves strong thrust with reduced lift fluctuations. The right panel illustrates the **maximum** displacement achieved in this experiment under parameterized motion, standard PPO training, and our proposed ACPPO-PID.

It is worth noting that for the *PPO without cost* algorithm, gait 1 shows relatively good displacement performance. To better understand this phenomenon, we further analyze the net forces and lift generated by the bio-inspired propulsors under the best-performing gait of each algorithm, computing the mean forward thrust (F_x^{mean}), mean lift (F_z^{mean}), and lift variance (F_z^{var}) over time. The results in Fig. 5 reveal clear quantitative differences. **ACPPO-PID (ours)** achieves a mean forward thrust of 0.93, which is approximately 27% higher than *CPPO-PID* and 86% higher than *BF*, while being only 24.7% lower than unconstrained *PPO*. More importantly, ACPPO-PID effectively suppresses destabilizing lift: its mean lift is 20.4% smaller than *PPO* and 61.9% smaller than *BF*, with only a slight increase of 10.3% compared with *CPPO-PID*. In terms of stability, ACPPO-PID shows the lowest lift variance among strong performers, reducing it by 55.9% relative to *PPO*, 49.2% relative to *BF*, and 16.7% relative to *CPPO-PID*. These results indicate that ACPPO-PID strikes the best trade-off in retaining strong propulsion while significantly reducing destabilizing forces. This allows the robot to maintain a more stable vertical position and thereby achieve greater net displacement.

A closer look at optimized gaits. Fig 6 compares limb motions between our method (top) and unconstrained *PPO* (bottom). Each cycle is divided into pre-paddle and paddle phases based on thrust generation, with key differences highlighted in calf movement (red) and hip movement (yellow). During the pre-paddle phase, our method adjusts the calf to maintain velocity direction parallel to calf orientation to reduce drag, while *PPO* fixes a small angle between calf and hip, positioning for greater extension and larger thrust-generating surface area in the subsequent phase. During the




Fig. 6. Comparison of limb motion patterns between our constrained method (ACPPPO-PID, top) and unconstrained PPO (bottom) during one locomotion cycle. Key differences are highlighted: calf movement during pre-paddle phase (red) and hip movement during paddle phase (yellow).

paddle phase, our method drives the hip backward directly, whereas PPO tilts forward then back. This creates a larger range of motion that enhances thrust generation but introduces significant additional lift.

Overall, PPO aggressively pursues thrust at the cost of lift-induced instability, while our constrained approach produces cleaner, more stable, and thrust-centric motion patterns that better achieve optimization objectives.

VI. CONCLUSION

This work demonstrates that explicit constraints on lateral forces and oscillatory amplitudes are essential for stable and efficient quadruped swimming. By formulating gait optimization as a constrained reinforcement learning problem, we enhance propulsion efficiency under strong fluid–structure interactions without sacrificing stability. To this end, we propose **ACPPPO-PID**, a safe RL framework for oscillatory propulsion that selectively widens exploration only when advantages are safe and positive and regulates updates with cycle-level performance. Experiments on both single actuators and quadruped robots confirm that ACPPPO-PID consistently outperforms strong baselines on the reward–cost frontier within limited hardware budgets. Current validation was limited to quiescent water, leaving performance under extreme hydrodynamic disturbances, such as strong currents or turbulence, as an open challenge. Future work will therefore focus on integrating online adaptation and domain randomization mechanisms to mitigate these environmental uncertainties and ensure robust deployment in open-water environments.

REFERENCES

- [1] Z. Wang, R. Lin, Z. Zhao, X. Chen, P. Guo, N. Yang, Z. Wang, and D. Fan, “Learn to flap: foil non-parametric path planning via deep reinforcement learning,” *Journal of Fluid Mechanics*, vol. 984, p. A9, 2024.
- [2] M. Triantafyllou, F. Hover, A. Techet, and D. Yue, “Review of hydrodynamic scaling laws in aquatic locomotion and fishlike swimming,” *Applied Mechanics Reviews*, vol. 58, no. 4, p. 226, 2005.
- [3] D. Xia, M. Lei, Z. Li, and Y. Shi, “A combined ib-lb method for predicting the hydrodynamics of bionic undulating fin thrusters,” *Ocean Engineering*, vol. 303, p. 117790, 2024.
- [4] S. Jung and R. Godoy-Diana, “bioinspired fluid–structure interaction,” *Bioinspiration & Biomimetics*, vol. 18, no. 3, p. 030401, 2023.
- [5] Y. Li, L. Chen, Y. Wang, and C. Ren, “Design and experimental evaluation of the novel undulatory propulsors for biomimetic underwater robots,” *Bioinspiration & Biomimetics*, vol. 16, no. 5, p. 056005, 2021.
- [6] A. Priovolos, E. Filippas, and K. Belibassakis, “A vortex-based method for improved flexible flapping-foil thruster performance,” *Engineering Analysis with Boundary Elements*, vol. 95, pp. 69–84, 2018.
- [7] D. Barrett, M. Triantafyllou, D. Yue, M. Grosenbaugh, and M. Wolfgang, “Drag reduction in fish-like locomotion,” *Journal of Fluid Mechanics*, vol. 392, pp. 183–212, 1999.

- [8] T. I. Fossen, “Handbook of marine craft hydrodynamics and motion control,” *John Wiley & Sons Ltd*, 2011.
- [9] J. E. Colgate and K. M. Lynch, “Mechanics and control of swimming: A review,” *IEEE journal of oceanic engineering*, vol. 29, no. 3, pp. 660–673, 2004.
- [10] G. Picardi, A. Astolfi, D. Chatziveangelou, J. Aguzzi, and M. Calisti, “Underwater legged robotics: Review and perspectives,” *Bioinspiration & Biomimetics*, vol. 18, no. 3, p. 031001, 2023.
- [11] B. Etkin and L. D. Reid, *Dynamics of flight: stability and control*. John Wiley & Sons, 1995.
- [12] F. Han, P. Guo, H. Chen, W. Li, J. Ren, N. Liu, N. Yang, and D. Fan, “Learn to swim: Data-driven lstm hydrodynamic model for quadruped robot gait optimization,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 10 261–10 267.
- [13] A. J. Ijspeert, “Central pattern generators for locomotion control in animals and robots: a review,” *Neural networks*, vol. 21, no. 4, pp. 642–653, 2008.
- [14] R. Tong, C. Qiu, Z. Wu, J. Wang, M. Tan, and J. Yu, “Na-cpg: A robust and stable rhythm generator for robot motion control,” *Biomimetic Intelligence and Robotics*, vol. 2, no. 4, p. 100075, 2022.
- [15] T. Kim, I. Hong, S. Im, S. Rho, M. Kim, Y. Roh, C. Kim, J. Park, D. Lim, D. Lee *et al.*, “Wing-strain-based flight control of flapping-wing drones through reinforcement learning,” *Nature Machine Intelligence*, vol. 6, no. 9, pp. 992–1005, 2024.
- [16] A. Ray, J. Achiam, and D. Amodei, “Benchmarking safe exploration in deep reinforcement learning,” *arXiv preprint arXiv:1910.01708*, vol. 7, no. 1, p. 2, 2019.
- [17] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *arXiv preprint arXiv:1707.06347*, 2017.
- [18] X. B. Peng, E. Coumans, T. Zhang, T.-W. Lee, J. Tan, and S. Levine, “Learning agile robotic locomotion skills by imitating animals,” *arXiv preprint arXiv:2004.00784*, 2020.
- [19] S. Choi, G. Ji, J. Park, H. Kim, J. Mun, J. H. Lee, and J. Hwangbo, “Learning quadrupedal locomotion on deformable terrain,” *Science Robotics*, vol. 8, no. 74, p. eade2256, 2023.
- [20] X. Cui, B. Sun, Y. Zhu, N. Yang, H. Zhang, W. Cui, D. Fan, and J. Wang, “Enhancing efficiency and propulsion in bio-mimetic robotic fish through end-to-end deep reinforcement learning,” *Physics of Fluids*, vol. 36, no. 3, 2024.
- [21] J. Achiam, D. Held, A. Tamar, and P. Abbeel, “Constrained policy optimization,” in *International conference on machine learning*. PMLR, 2017, pp. 22–31.
- [22] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, “Trust region policy optimization,” in *International conference on machine learning*. PMLR, 2015, pp. 1889–1897.
- [23] Y. Zhang, Q. Vuong, and K. Ross, “First order constrained optimization in policy space,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 15 338–15 349, 2020.
- [24] T.-Y. Yang, J. Rosca, K. Narasimhan, and P. J. Ramadge, “Projection-based constrained policy optimization,” *arXiv preprint arXiv:2010.03152*, 2020.
- [25] L. Yang, J. Ji, J. Dai, Y. Zhang, P. Li, and G. Pan, “Cup: A conservative update policy algorithm for safe reinforcement learning,” *arXiv preprint arXiv:2202.07565*, 2022.
- [26] C. Tessler, D. J. Mankowitz, and S. Mannor, “Reward constrained policy optimization,” *arXiv preprint arXiv:1805.11074*, 2018.
- [27] A. Stooke, J. Achiam, and P. Abbeel, “Responsive safety in reinforcement learning by pid lagrangian methods,” in *International Conference on Machine Learning*. PMLR, 2020, pp. 9133–9143.
- [28] M. Willis, “Proportional-integral-derivative control,” *Dept. of Chemical and Process Engineering University of Newcastle*, vol. 6, 1999.
- [29] R. S. Sutton, A. G. Barto *et al.*, *Introduction to reinforcement learning*. MIT press Cambridge, 1998, vol. 135.
- [30] Q. Yu, Z. Zhang, R. Zhu, Y. Yuan, X. Zuo, Y. Yue, W. Dai, T. Fan, G. Liu, L. Liu *et al.*, “Dapo: An open-source llm reinforcement learning system at scale,” *arXiv preprint arXiv:2503.14476*, 2025.
- [31] Y. Zhao, Y. Liu, J. Liu, J. Chen, X. Wu, Y. Hao, T. Lv, S. Huang, L. Cui, Q. Ye *et al.*, “Geometric-mean policy optimization,” *arXiv preprint arXiv:2507.20673*, 2025.