

# S-Cheetah: A Novel Quadrupedal Robot with a 3-DOF Active Spine Learning Agile Locomotion

Zimu Li and Weibang Bai\*

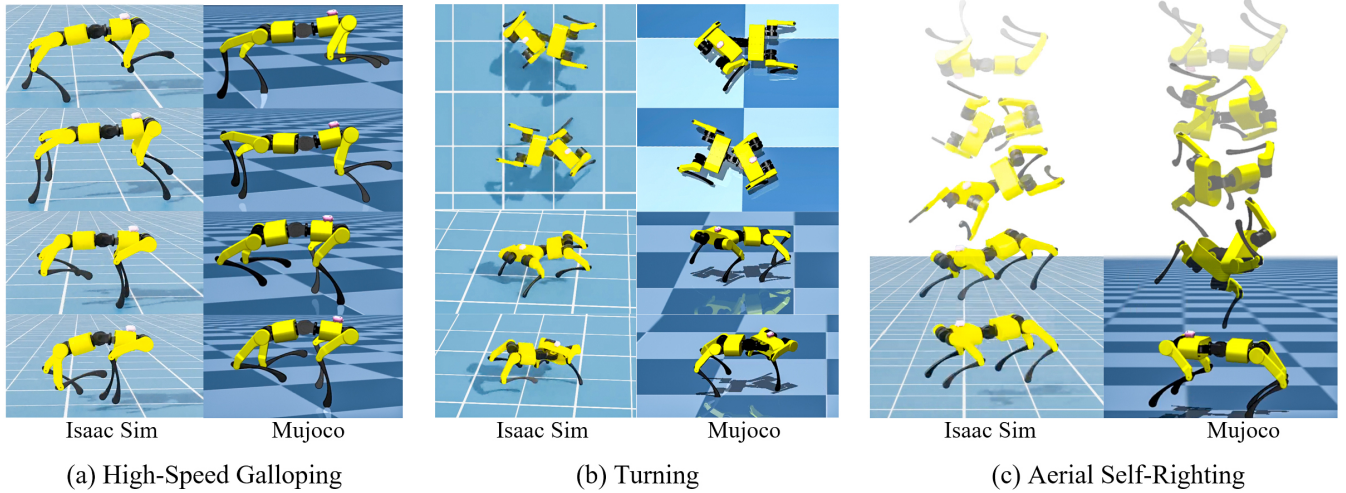


Fig. 1: *S-Cheetah* 🐆: A 3-DOF *Spine-driven Cheetah-inspired Quadrupedal Robot for Agile Locomotion*. Our proposed S-Cheetah demonstrates agile locomotion capabilities across diverse scenarios. (a) High-speed rotary G2 galloping. (b) Maneuverable in-place turning and moving turning. (c) Emergent feline-inspired aerial self-righting.

**Abstract**—The biological spine of quadrupeds enables sagittal flexion/extension, lateral bending, and axial rotation, playing a crucial role in highly agile and dexterous locomotion. While numerous studies have integrated active spinal joints into quadrupedal robots to enhance agility, most designs simplify control complexity by reducing spinal degrees of freedom (DOF), failing to achieve the spatial tri-axial rotation characteristic of biological spines. Consequently, replicating a multi-DOF biomimetic spine and effectively leveraging it to empower the agile locomotion of quadrupedal robots remains a significant research challenge. In this study, we present S-Cheetah, a quadrupedal robot featuring a 3-DOF bio-inspired serial active spine capable of biomimetic spatial tri-axial rotation. To empower the robot to fully utilize this active spine, we developed a specialized reinforcement learning framework to actively promote the engagement of the introduced spine and maximize the robot’s locomotive capabilities by integrating an acceleration curriculum learning strategy with tailored reward functions, such as a gallop gait reward, a spine undulation reward, and a spine steering reward. Experimental results demonstrate that S-Cheetah can achieve a peak speed of 6.9 m/s using the rotary G2 gallop gait and an in-place turning rate of 7.2 rad/s. Besides, the system exhibits an emergent, feline-inspired aerial self-righting capability, allowing it to land stably on four feet from arbitrary orientations during free fall. Finally, through extensive evaluations across diverse locomotion tasks, we prove that the introduction of the proposed 3-DOF spine comprehensively enhances the locomotive agility of quadrupedal robots. Project website: [himmy-robotics.github.io/scheetah](https://himmy-robotics.github.io/scheetah)

## I. INTRODUCTION

Through hundreds of millions of years of evolution, terrestrial quadrupeds have developed a sophisticated synergy between their limbs and spinal column, endowing them with extraordinary terrain adaptability and locomotor agility [1]. This is manifested in their diverse gaits, from the walk and trot to the high-speed gallop [2]. Throughout agile locomotion, the spine serves not merely as a structural support but as a central engine for power modulation. Particularly during galloping, the rhythmic flexion and extension of the spinal joints significantly amplify stride length, providing substantial gains in dynamic performance [3].

Inspired by these biological archetypes, the field of quadrupedal robotics has made remarkable progress. The integration of high-torque-density actuators [4] and reinforcement learning-based control [5] has enabled modern quadrupedal robots to achieve unprecedented agility, including complex parkour maneuvers in challenging environments [6], [7]. Notably, the “Black Panther” quadrupedal robot achieved a top speed of 10.3 m/s [8], and the “White Rhino” set a Guinness World Record for the fastest 100 m by a quadrupedal robot at 16.33 seconds [9]. However, in the pursuit of control simplicity and structural robustness, these high-performance platforms universally adopt a rigid-trunk design, omitting the degrees of freedom (DOF) provided by a biological spine. Consequently, even at peak velocities, they predominantly rely on the highly stable yet dynamically conservative trot gait rather than the gallop, which is known

for its explosive power and speed in nature [2], fundamentally limiting further improvements in robotic dynamic performance.

To bridge this gap, numerous studies have explored integrating articulated spinal joints into quadrupedal robots. Bhattacharya et al. [10] designed a quadrupedal robot with a 2-DOF pitch-actuated spine, achieving a running speed of 2.21 m/s in simulation via reinforcement learning. Caporale et al. [11] developed a quadrupedal robot featuring a single roll-DOF spine, experimentally demonstrating that spinal twisting during locomotion reduces energy consumption. Similarly, Qian et al. [12] designed a quadruped with a roll-DOF spine, showing through experiments that the introduction of spinal rolling reduces the turning radius while enhancing path-tracking capability and adaptability to complex terrain. Furthermore, Yoneda et al. [13] utilized reinforcement learning to enable a quadrupedal robot with a single pitch-DOF spine to perform wall-climbing maneuvers. The “Yat-sen Lion II” robot presented by Yang et al. [14] features a 2-DOF spine (pitch and yaw), demonstrating that spinal integration results in a smaller turning radius and more agile trajectory tracking performance. Some research has explored higher DOFs. Wang et al. [15] employed a symmetric five-bar mechanism to construct a 4-DOF spine capable of twisting, extension, contraction, and rotation. However, it lacks the extensive range of motion found in biological quadrupeds, and the spine functions as a passive spring rather than an actively controlled component. Wu et al. [16] introduced a novel spinal structure where each module possesses two DOFs, allowing for serially connected modules to increase the total DOF, and demonstrated aerial self-righting during free-fall through spinal articulation. However, since each module is constrained to planar pitch and yaw motions, the assembled spine still lacks a roll DOF. Moreover, the complex multi-module spinal design coupled with the simplified leg structure limits the robot’s ability to perform other complex locomotion tasks. Regarding speed performance, the robot developed by Hu et al. [17] achieved a maximum galloping speed of 4.9 m/s in simulation, yet the pitch spinal joint exhibited only negligible angular variation. The “Yat-sen Lion” robot introduced by Li et al. [18] attained a top bounding speed of 4.0 m/s in simulation. Its two pitch spinal joints (front and rear) reached a maximum angle of 0.21 rad ( $\approx 12.0^\circ$ )—equivalent to a single joint angle of 0.42 rad ( $\approx 24.1^\circ$ ). However, the spinal angles reached during flexion and extension were identical, which is inconsistent with the asymmetric spinal motion patterns observed in biological quadrupeds during running. For instance, a galloping cheetah’s spine achieves a maximum flexion angle of  $84^\circ$  versus an extension angle of only  $44^\circ$  [19].

The unique biological structure of the quadrupedal spine enables spatial tri-axial rotation, which plays a crucial role not only in enhancing locomotive agility but also in enabling remarkable behaviors such as feline aerial self-righting during falls [20]. However, current research on spinal quadrupedal robots rarely achieves this bio-inspired

mode of tri-axial rotation, and although Wu et al. [16] demonstrated spinal self-righting, no existing work has realized both agile locomotion and aerial self-righting within a unified framework. Furthermore, regarding control methods, whether using reinforcement learning [10], [13], [15] or Model Predictive Control (MPC) [14], [16], [17], [18], a major challenge remains: how to effectively utilize the spine to holistically enhance the locomotive agility of quadrupedal robots.

To address these challenges, we propose S-Cheetah: a novel quadrupedal robot featuring a bio-inspired serial 3-DOF active spine. This spine is capable of spatial tri-axial rotation, enabling it to accurately replicate the spinal kinematics of biological quadrupeds. Furthermore, we developed a specialized reinforcement learning framework with tailored reward functions that effectively synergizes spinal articulation with limb movements, thereby achieving superior locomotive agility.

The main contributions of our work are summarized as follows:

- 1) We develop a novel quadrupedal robot integrated with a 3-DOF bio-inspired serial active spine, capable of accurately replicating the complex tri-axial spinal rotations observed in biological quadrupedal locomotion.
- 2) We propose a reinforcement learning framework with tailored reward functions that effectively coordinates spinal and limb motions, enabling the robot to achieve high-speed galloping, maneuverable turning, precise path-tracking, and emergent feline-inspired aerial self-righting.
- 3) We conduct comprehensive evaluations across multiple highly agile locomotion tasks in both Isaac Sim and MuJoCo simulation environments, demonstrating that the 3-DOF active spine significantly enhances the robot’s swiftness, maneuverability, precision, and stability.

The remainder of this paper is organized as follows. Section II presents the system design of S-Cheetah. Section III introduces the reinforcement learning method used for control. In Section IV, simulation experiments verify the reliability of the proposed method, followed by a conclusion of the paper in Section V.

## II. SYSTEM DESIGN

### A. Overview of Structure

As shown in Fig. 2, the quadrupedal robot S-Cheetah has a total of 15 DOFs, including 3 DOFs for each leg and 3 DOFs for the spine. The rotation axis of each joint is indicated by the red arrows in the figure, while the corresponding rotational motion is illustrated by the green arrows.

The trunk of S-Cheetah consists of two parts, namely the front-body-link and the rear-body-link, which are connected via a spine module. Each leg is composed of a hip-link, a thigh-link, and a calf-link. The robot has a total mass of 20 kg, an overall body length of 625 mm, a width of 380 mm, and a leg length of 580 mm, with a standing height




Fig. 2: Structural composition and joint rotation representation of S-Cheetah.




Fig. 3: Design of S-Cheetah's spine joints and their correspondence to a quadrupedal spine.

of approximately 440 mm. The peak torque of the leg joint motors is  $33.5 \text{ N} \cdot \text{m}$ . To more effectively power the robot's locomotion, the spine joint motors are designed with a higher peak torque of  $50 \text{ N} \cdot \text{m}$ .

### B. Design of Spine Joint

As shown in Fig. 3, the spine module consists of three serially connected joints, arranged from front to rear as the yaw-spine-link, pitch-spine-link, and roll-spine-link. The angular ranges of the yaw, pitch, and roll spine joints are  $\pm 1 \text{ rad}$  ( $\approx \pm 57.3^\circ$ ),  $\pm 1.5 \text{ rad}$  ( $\approx \pm 85.9^\circ$ ), and  $\pm 2 \text{ rad}$  ( $\approx \pm 114.6^\circ$ ), respectively. The ordering of the three serial spine joints is determined based on the following considerations:

- 1) The yaw-spine-link is placed at the foremost position, allowing the front half of the robot to yaw independently. This mimics the gaze-oriented steering behavior observed in animals during prey pursuit, deflecting the forelegs' workspace relative to the hind legs to guide the body naturally into a turn.
- 2) The pitch-spine-link is positioned at the middle of the spine module. Placing it at the mechanical center between the fore and hind hip joints maximizes the stride-length gain from spinal flexion and extension, replicating the bow-like elastic energy cycle of the cheetah's lumbar spine. An off-center placement would create unbalanced moment arms, reducing the effectiveness of spinal energy storage.

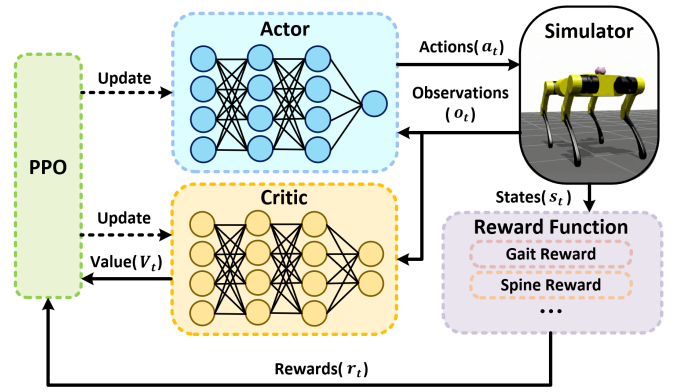


Fig. 4: Basic framework of the reinforcement learning method used for training.

- 3) The roll-spine-link is located at the rear end of the spine module, near the pelvis and hind legs. This allows the posterior trunk to twist relative to the anterior trunk, emulating the independent pelvic adjustments animals perform on uneven terrain and the lateral roll of the hind legs to counteract centrifugal forces during high-speed cornering.

The introduction of these three serial spine joints enables comprehensive replication of the dexterous spinal motions of quadrupeds in roll, pitch, and yaw. Moreover, the spine configuration characterized by “guidance prioritized at the front, power centralized in the middle, and compliance at the rear” allows S-Cheetah to exhibit locomotion dynamics that more closely resemble those of biological quadrupeds.

## III. REINFORCEMENT LEARNING-BASED CONTROL FOR AGILE LOCOMOTION

### A. RL Framework

As shown in Fig. 4, the design of the reinforcement learning framework is based on the RSL-RL architecture [21]. Two networks, namely the actor and the critic, are trained simultaneously: the actor generates actions, while the critic estimates the value function. Proximal Policy Optimization (PPO) [22] is adopted as the learning algorithm to update the parameters of both the actor and critic networks.

During control execution, only the actor network is employed. The actor outputs scaled target joint angles for 15 joints (12 leg joints and 3 spine joints) at a frequency of 50 Hz. These target joint angles are then fed into a simple PD position controller, which computes the corresponding joint torques to drive the robot's motion.

### B. Curriculum Learning Design

Training agile locomotion behaviors such as galloping and turning directly with full-range velocity commands is challenging, as the agent struggles to explore effectively under high-speed targets from the outset. We address this with a two-fold curriculum learning strategy. First, we apply an adaptive velocity range curriculum: the commanded linear and angular velocity ranges are initialized to 10% of their

full target ranges and progressively expanded as training performance improves. Specifically, when the mean episodic tracking reward exceeds a predefined threshold (80% for linear velocity, 60% for angular velocity), the range is incrementally increased by  $\Delta = 0.1$  per training iteration until the full range is reached. The linear and angular velocity curricula operate independently, allowing the agent to master forward locomotion and yaw control at their own pace.

Second, we introduce an acceleration-constrained command generator to avoid abrupt velocity jumps that the robot cannot physically track. Rather than issuing step-wise target velocities, the commanded velocity  $\mathbf{v}_t$  is smoothly ramped toward the target  $\mathbf{v}^*$  at each simulation step:

$$\mathbf{v}_{t+1} = \mathbf{v}_t + \text{clamp}(\mathbf{v}^* - \mathbf{v}_t, -\mathbf{a}_{\max}\Delta t, \mathbf{a}_{\max}\Delta t), \quad (1)$$

where  $\Delta t$  is the control time step and  $\mathbf{a}_{\max}$  is sampled from a predefined range. For the turning task, the angular velocity command is further delayed by a random interval after the linear velocity command, enabling the robot to first reach a stable forward speed before initiating a turn. This combined curriculum significantly improves training stability and final performance.

### C. Main Reward Function Design

1) *Gait Reward*: Gallop is the primary gait exhibited by quadrupeds during maximal-speed locomotion. A previous study [23] provides a detailed classification of galloping gaits: based on limb-ground contact sequences, gallop is divided into transverse and rotary; based on the number of aerial phases, it is categorized into G0, GG, GE, and G2. Agile quadrupeds like the cheetah typically adopt the rotary G2 gallop, characterized by two aerial phases, as it optimally balances velocity, stability, and energy efficiency. Accordingly, we design a gait reward function to guide the policy toward learning this specific gait for high-speed locomotion.

Fig. 5 illustrates the locomotion performance and gait diagrams of the rotary G2 gallop. The diagrams depict the temporal sequence of limb-ground contacts (where LH, RH, LF, and RF denote the left/right hind/forelimbs, respectively), with gray segments indicating aerial phases where no limbs touch the ground.

The core idea of our gait reward design is to specify only the minimal footfall pattern distinguishing rotary gallop from other gaits, without prescribing the full gait cycle, duty factor, or stride frequency. In gallop, ipsilateral limb pairs exhibit a characteristic temporal asymmetry: the leading leg contacts the ground slightly before the trailing leg with a non-zero phase offset. This distinguishes gallop from bound, where ipsilateral limbs land simultaneously. To capture this, we define a phase-offset reward based on a Gaussian kernel. For each ipsilateral pair (front and rear), we compute the temporal phase offset between their ground-contact events. To account for the periodic nature of the gait cycle, we calculate the shortest phase distance  $\Delta\delta$  between the measured

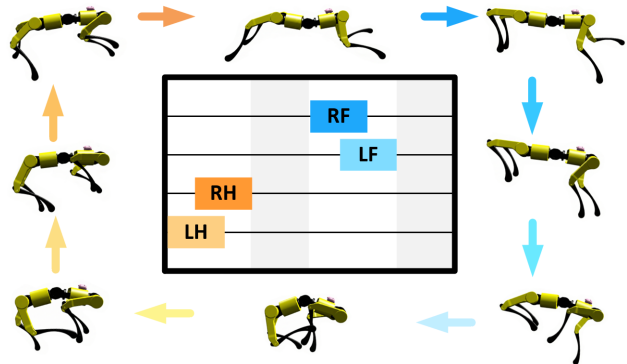


Fig. 5: Rotary G2 gallop performance and its gait diagrams.

offset and the target offset  $\delta^*$ . The reward is formulated as:

$$r_{\text{pair}} = \exp\left(-\frac{\Delta\delta^2}{\sigma^2}\right), \quad (2)$$

where  $\sigma$  controls the tolerance around the target offset.

The total gait reward is the product of the front and rear pair rewards:  $r_{\text{gait}} = r_{\text{front}} \cdot r_{\text{rear}}$ . Crucially, we impose no explicit constraint on the phase relationship between the front and rear limb groups, nor on the duration of aerial phases. By setting different target offsets for the front ( $\delta_f^*$ ) and rear ( $\delta_r^*$ ) pairs, the reward provides sufficient guidance to prevent bounding ( $\delta^* = 0$ ) while preserving flexibility. This deliberate under-specification, combined with auxiliary rewards like air time and velocity tracking, allows the policy to naturally discover the full rotary G2 gallop—including its two distinct aerial phases and characteristic circular footfall sequence (LH→RH→RF→LF)—as an emergent optimal solution.

2) *Spine Undulation Reward*: In prior work [10], no dedicated reward was designed to actively drive spinal joint motion; instead, the spine was expected to learn useful movements implicitly through velocity tracking rewards and power penalties. While this passive approach does produce some spinal oscillation during locomotion, it is highly susceptible to reward hacking; the policy tends to converge to minimal or irregular spinal movements that satisfy the tracking objective without leveraging the spine to enhance propulsion. In nature, agile quadrupeds utilize a rhythmic, wave-like spinal motion known as undulation to maximize stride length and generate explosive forward thrust. To induce this crucial biomechanical feature, we design an explicit spine undulation reward for the pitch spinal joint. This reward encourages large-amplitude, phase-coordinated spinal flexion-extension to actively power high-speed galloping.

The reward consists of three components. First, a phase coordination term evaluates whether the spine moves in the biomechanically correct direction relative to the limbs. We compute an effective leg velocity  $\bar{v}_{\text{leg}}$  by averaging the rear thigh velocities with the sign-inverted front thigh velocities, so that a positive value consistently indicates a phase requiring spinal extension. The phase score is then:

$$r_{\text{phase}} = \tanh(\dot{\theta}_{\text{pitch}} \cdot \bar{v}_{\text{leg}} \cdot \alpha), \quad (3)$$

where  $\dot{\theta}_{\text{pitch}}$  is the spine angular velocity and  $\alpha$  is a sensitivity coefficient.

Second, an amplitude incentive encourages the spine to reach a target deflection angle. Based on the sign of  $\bar{v}_{\text{leg}}$ , the expected spine state  $s \in \{-1, +1\}$  is determined (flexion or extension), and the amplitude reward is computed as:

$$r_{\text{amp}} = s \cdot \theta_{\text{pitch}}. \quad (4)$$

During quadrupedal locomotion, there is a substantial disparity in the spinal angular limits between flexion and extension [19]. Typically, the maximum angle achieved during flexion significantly exceeds that of extension. Consequently, we define an expected threshold  $\theta_{\text{th}}$  (which denotes either  $\theta_{\text{flex}}$  or  $\theta_{\text{ext}}$  depending on the spine state) to faithfully simulate the spinal dynamics. The effective amplitude reward is then clamped as  $\bar{r}_{\text{amp}} = \text{clamp}(r_{\text{amp}}, 0, \theta_{\text{th}})$ .

Third, an excess penalty prevents over-rotation beyond the desired range:

$$p = w_e (\max(0, r_{\text{amp}} - \theta_{\text{th}}))^2, \quad (5)$$

where  $w_e$  is the excess penalty.

The final reward combines these as:

$$R_{\text{spine\_undulation}} = r_{\text{phase}} \cdot (1 + w_b \cdot \bar{r}_{\text{amp}}) - p, \quad (6)$$

where  $w_b$  is the amplitude boost weight and  $\bar{r}_{\text{amp}}$  is the clamped amplitude reward. This design ensures that the spine is rewarded for moving in the correct phase with sufficient amplitude, while being penalized for exceeding safe mechanical limits.

3) *Spine Steering Reward*: Conventional rigid-trunk quadrupedal robots can only achieve turning through differential leg velocities between the left and right sides, making it inherently difficult to decouple forward locomotion from yaw maneuvers. The introduction of a yaw-axis spinal joint addresses this limitation, functioning analogously to the steering mechanism of a vehicle that directs the front axle. To encourage the policy to actively utilize this joint for turning, we design a spine steering alignment reward.

The reward incentivizes the yaw spine joint to deflect in the direction consistent with the commanded angular velocity. To ensure that the spine is only activated during intentional turning maneuvers, we incorporate a dead-zone threshold  $\omega_{\text{th}}$  into a piecewise reward function:

$$R_{\text{spine\_steering}} = \begin{cases} 0, & \text{if } |\omega_z^*| < \omega_{\text{th}} \\ \tanh\left(\frac{-\theta_{\text{yaw}}}{\omega_z^*} \cdot k\right), & \text{otherwise} \end{cases}, \quad (7)$$

where  $k$  is a scaling factor controlling how quickly the reward saturates. When the yaw rate command is within the dead-zone ( $|\omega_z^*| < \omega_{\text{th}}$ ), the reward is strictly zeroed out, effectively preventing any spurious spine deflection during straight-line locomotion. When an active turn is commanded ( $|\omega_z^*| \geq \omega_{\text{th}}$ ), the reward evaluates the alignment. Crucially, the formulation inherently encapsulates the sign of the command in the denominator  $\omega_z^*$ . This ensures that a positive reward is given when the spine deflects in the correct steering direction, while an incorrect deflection yields a negative penalty, providing a clear directional learning signal.

## IV. EXPERIMENTS

To validate the effectiveness of the proposed spinal quadrupedal robot and reward design, we conducted extensive experiments across two physics engines. Our training framework, built upon the open-source codebase [24], leverages NVIDIA Isaac Sim to enable massively parallel, GPU-accelerated reinforcement learning. The learned policy was first evaluated in Isaac Sim and then directly transferred to MuJoCo—a simulator favored in the legged robotics community for its high-fidelity contact dynamics and minimal sim-to-real gap. This seamless cross-engine transfer, achieved without fine-tuning or domain adaptation, underscores the policy’s robustness and provides encouraging evidence for future physical deployment.

To isolate the contribution of the articulated spine, we compared two configurations: (1) the full robot model with a 3-DOF active spine, and (2) a rigid-trunk baseline with locked spinal joints. Both configurations shared identical hyperparameters, observation/action spaces, and reward functions (excluding spine-specific terms). Furthermore, both were trained for a consistent number of iterations to ensure a fair and rigorous comparison.

### A. Spine-Empowered High-Speed Galloping

The rotary G2 gallop, the fastest gait employed by agile quadrupeds like cheetahs, relies on coordinated spinal flexion and extension to augment limb propulsion. As illustrated in Fig. 1(a), our proposed framework successfully realizes this complex asymmetric gait, effectively leveraging spinal dynamics to enhance locomotor performance.

To evaluate speed limits, we applied high-velocity tracking commands across both simulators (Fig. 6). In Isaac Sim, the active-spine robot achieved a stable speed of 6.9 m/s, representing a 15% improvement over the 6.0 m/s rigid-trunk baseline. In MuJoCo, the performance gain was even more pronounced: the active-spine robot reached 6.0 m/s, representing a 25% enhancement over the baseline’s 4.8 m/s. Notably, both results significantly exceed the current state-of-the-art (SOTA) for simulated spined quadrupeds, which stands at 4.9 m/s [17].

Kinematic analysis of the spinal joints in Fig. 7 reveals that the pitch joint exhibits substantial asymmetric variations, with angular ranges of  $[-0.5, 0.25]$  rad in Isaac Sim and  $[-0.6, 0.3]$  rad in MuJoCo. This behavior aligns with biological principles where spinal flexion magnitude is approximately double that of extension during galloping. Meanwhile, the roll joint maintains minimal variation, ensuring lateral stability. This bio-inspired asymmetry contrasts sharply with existing MPC-based methods [18], which prescribe symmetric spinal motion (e.g.,  $[-0.42, 0.42]$  rad), a simplification that contradicts the natural kinematic characteristics of rapid galloping.

### B. Spine-Assisted Agile Turning

1) *Turning Capability Evaluation*: To evaluate the turning enhancements enabled by the articulated spine, we conducted angular velocity tracking experiments in both simulators. We

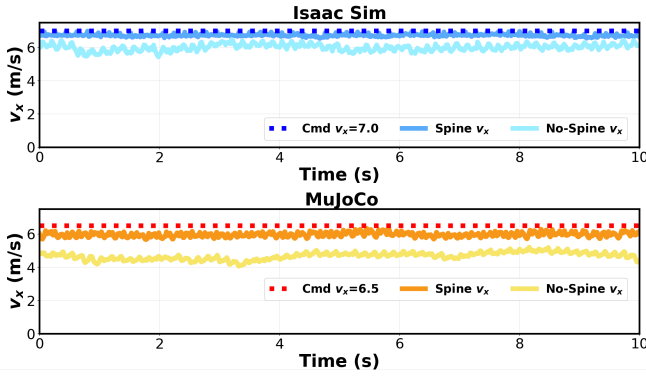


Fig. 6: The  $v_x$  tracking performance during galloping.

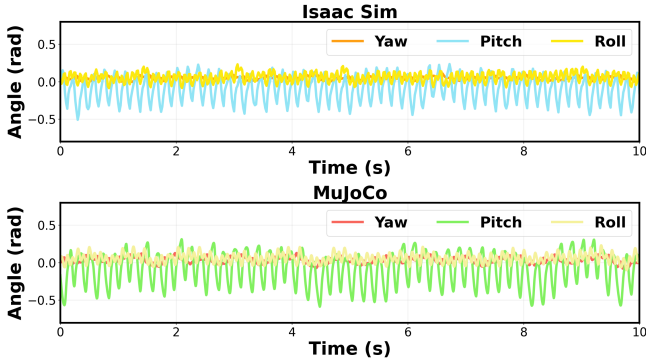


Fig. 7: Angular variations of the three spine joints during galloping.

measured the maximum yaw rate ( $\omega_z$ ) for the active-spine and rigid-trunk configurations under two conditions: in-place turning and moving turning (tracking a constant forward velocity of 1 m/s), as depicted in Fig. 1(b).

As shown in Fig. 8, the active-spine configuration consistently and significantly outperformed the rigid-trunk baseline. During in-place turning, the active-spine robot achieved a stable maximum  $\omega_z$  of 7.2 rad/s in Isaac Sim, representing a 16.1% improvement over the rigid baseline, and 5.3 rad/s in MuJoCo, a 60.1% increase compared to its rigid counterpart. The performance gap further widened during moving turning tasks: the active-spine robot reached 7.8 rad/s in Isaac Sim, marking a 22.9% gain over the rigid baseline, and 6.6 rad/s in MuJoCo, an impressive 83.3% enhancement relative to the rigid-trunk model. Notably, these yaw rates substantially surpass the 3 rad/s state-of-the-art previously reported for quadrupeds with yaw-actuated spines [18].

Kinematic analysis of the spinal joints in Fig. 9 reveals the underlying mechanisms driving this agility. During high-speed turns, the spinal yaw joint operates near its mechanical limits to maximize rotational capability. Concurrently, the spinal roll joint actively adjusts to execute a banking maneuver—leaning into the curve. This bio-inspired compensation effectively counteracts centrifugal forces, significantly enhancing the robot’s dynamic stability during rapid directional changes.

2) *Path Tracking Performance:* To further investigate how the 3-DOF spine enhances maneuverability, we conducted a “figure-8” path-tracking task in MuJoCo. The trajectory

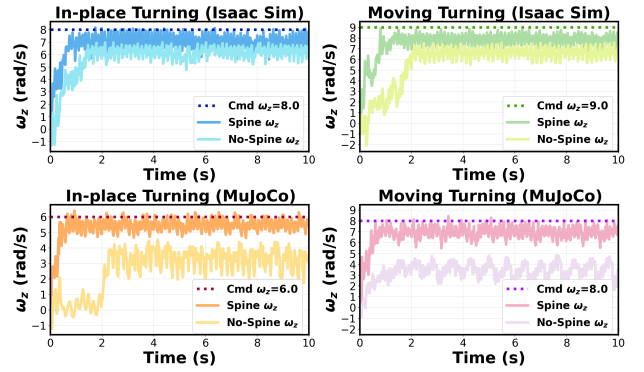


Fig. 8: The  $\omega_z$  tracking performance during turning.

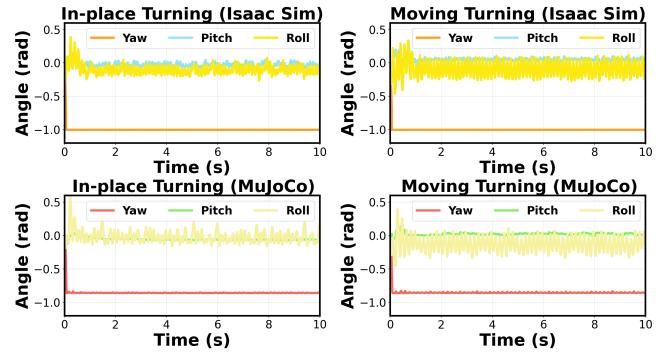


Fig. 9: Angular variations of the three spine joints during turning.

comprises two circular arcs ( $R = 1$  m) connected by intersecting straight segments ( $L = 4$  m). During the task, the robot maintains a constant commanded forward velocity  $v_x$ , while a proportional (P) controller regulates steering by translating heading errors into target yaw rates  $\omega_z$  for the policy to track.

As illustrated by the spatial trajectories in Fig. 10(a), the active-spine robot tracks the reference path more precisely than the rigid-trunk baseline, leveraging its mid-body DOF for fluid steering. Taking the  $v_x = 2.0$  m/s case as a representative example, quantitative analysis in Fig. 10(b) demonstrates that the active-spine configuration generates a higher transient yaw rate when entering curves, enabling more rapid heading adjustments than the rigid-trunk baseline. The spinal kinematic profiles in Fig. 10(c) further reveal this coordination: as the robot transitions into turns, the spinal yaw angle increases to tighten the turning radius, while the roll angle simultaneously adjusts to provide banking support.

Table I summarizes the performance metrics across various velocities. The active-spine robot consistently outperforms the rigid-trunk counterpart, achieving shorter completion times while maintaining superior accuracy in both path following and velocity tracking. Notably, these consistent improvements are achieved under a straightforward tracking framework with identical velocity commands, rather than using a specialized planner trained to push the physical limits of either platform. While beyond the scope of this work, we anticipate that a dedicated time-optimal policy could further magnify this performance gap, as the active spine provides

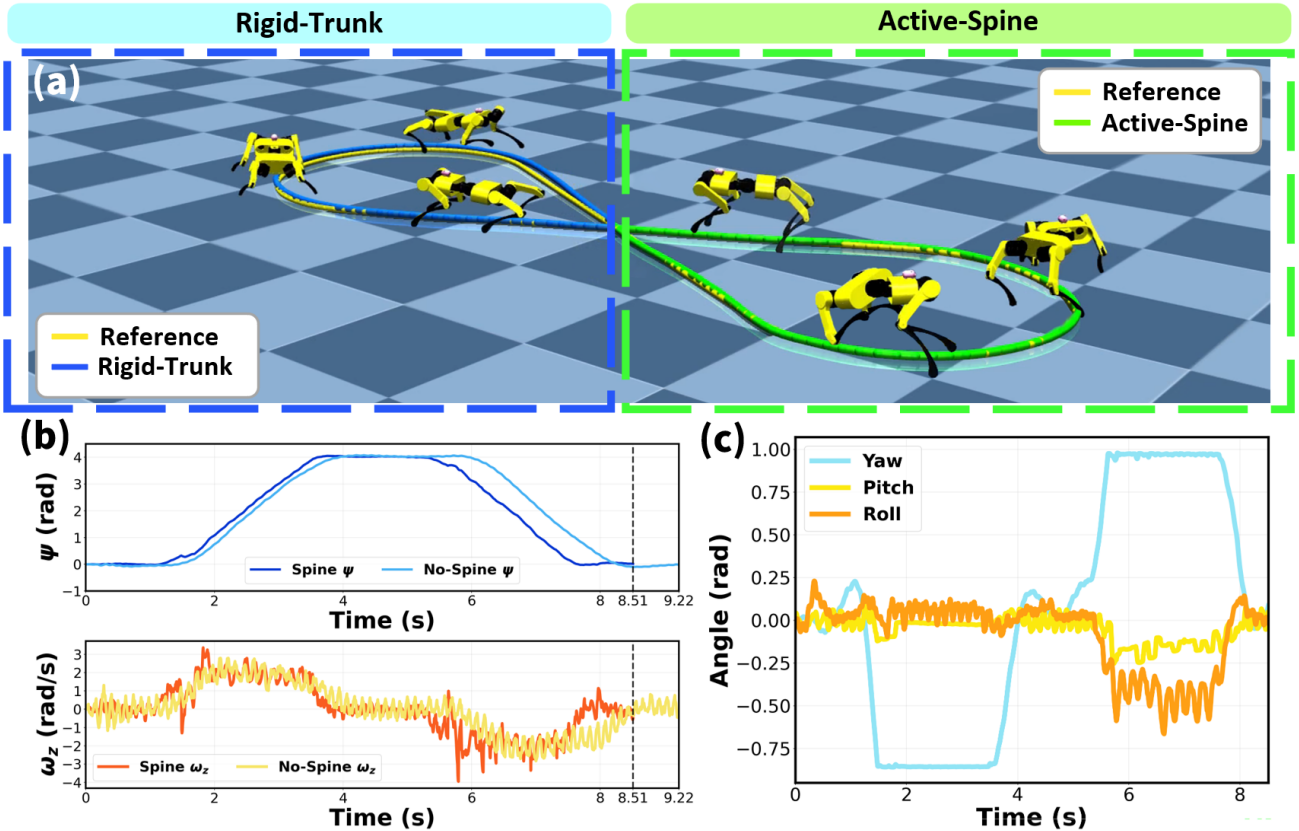


Fig. 10: Evaluation of the “figure-8” path-tracking performance. (a) Spatial trajectories of the active-spine and rigid-trunk robots. (b) Variations in yaw angle  $\psi$  and yaw rate  $\omega_z$  for both configurations. (c) Angular variations of the three spinal joints during path-tracking.

greater mechanical headroom for extreme maneuvers.

TABLE I: Path-tracking performance of rigid-trunk and active-spine robots at various velocities  $v_x$ .

$v_x$ (m/s)	Total Time (s)		Average Path Tracking Error (m)		Average Velocity Tracking Error (m/s)	
	Rigid-Trunk	Active-Spine	Rigid-Trunk	Active-Spine	Rigid-Trunk	Active-Spine
1.0	16.995	14.695	0.039	0.036	0.118	0.118
1.5	11.655	10.535	0.037	0.035	0.187	0.113
2.0	9.220	8.510	0.044	0.034	0.251	0.139
2.5	8.000	7.565	0.059	0.046	0.457	0.393

### C. Spine-Driven Aerial Self-Righting

During the training process, we observed an emergent capability in the active-spine robot: a markedly higher success rate in achieving stable four-point landings from random initial poses, a task where the rigid-trunk counterpart frequently suffered body-ground collisions. To systematically evaluate this, we designed an aerial drop experiment from a height of 2 m under three extreme initial orientations: (pitch =  $90^\circ$ ), (roll =  $135^\circ$ , pitch =  $45^\circ$ ), and a complete inversion (roll =  $180^\circ$ ).

As illustrated in Fig. 1(c), the active-spine robot leverages its 3-DOF spinal joints to rapidly modulate its body orientation during descent. Taking the complete inversion ( $\phi = 180^\circ$ ) as a representative case, the robot coordinates its spinal rotations to reorient its chassis mid-air, ensuring its feet are optimally positioned for impact, closely mimicking the feline self-righting reflex. Quantitative analysis in Fig. 11 confirms that while the rigid-trunk robot fails to significantly

modulate its roll angle  $\phi$  due to a lack of internal momentum manipulation, the active-spine robot generates substantial roll rate  $\omega_x$  to correct its posture, driven by the coordinated joint-angle trajectories shown in Fig. 12. Conversely, the rigid-trunk model often impacts the ground on its back, relying solely on post-landing limb recovery.

Table II summarizes the landing success rates across both simulators. Under the most challenging inverted orientation ( $\phi = 180^\circ$ ), the rigid-trunk robot exhibits a near-zero success rate. In stark contrast, the active-spine robot achieves near-100% success in Isaac Sim and 88% in MuJoCo, demonstrating that the articulated spine is indispensable for robust aerial autonomy and landing stability.

TABLE II: Success rates (%) under different initial orientations in two simulation environments.

Init. Orient	Robot Cfg			
	Rigid-trunk (Isaac Sim)	Active-spine (Isaac Sim)	Rigid-trunk (MuJoCo)	Active-spine (MuJoCo)
$p = 90^\circ$	60	> <b>98</b>	5	<b>96</b>
$r = 135^\circ, p = 45^\circ$	25	> <b>98</b>	28	<b>90</b>
$r = 180^\circ$	< 1	> <b>98</b>	< 1	<b>88</b>

## V. CONCLUSION

In this study, we introduce S-Cheetah, a quadrupedal robot featuring a bio-inspired serial 3-DOF active spine, which can effectively replicate the tri-axial rotation of biological quadrupeds. By leveraging a specialized reinforcement learning framework with tailored reward functions, the robot

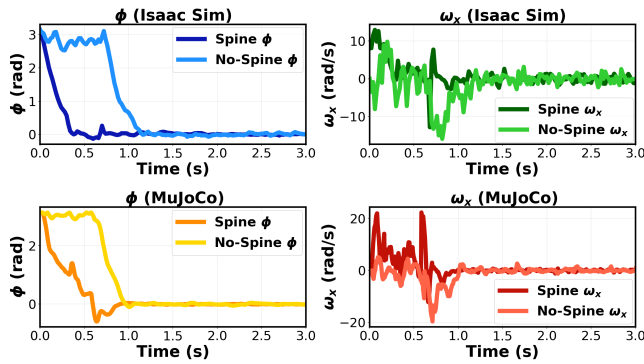


Fig. 11: Variations in roll angle and roll rate during aerial self-righting.

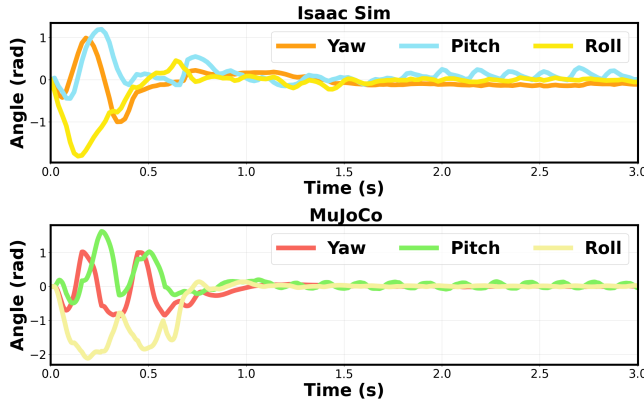


Fig. 12: Angular variations of the three spine joints during the aerial self-righting.

achieves agile locomotion, including high-speed galloping, maneuverable turning, precise path-tracking, and emergent feline-inspired aerial self-righting. Through comprehensive cross-validation in Isaac Sim and MuJoCo, we established a highly robust policy with a minimal sim-to-sim gap. In these simulated environments, S-Cheetah achieved a peak running speed of 6.9 m/s using the rotary G2 gallop gait and an in-place turning rate of 7.2 rad/s, substantially outperforming rigid-trunk baselines and current state-of-the-art simulated spinal quadrupedal robots. Crucially, our analysis confirms that the integrated 3-DOF spine significantly enhances the robot’s swiftness, maneuverability, precision, and stability.

Although this study currently focuses on the design and simulation phases prior to hardware deployment, it provides strong computational evidence that a bio-inspired serial 3-DOF active spine can empower quadrupedal robots with a level of locomotion that closely rivals the agility of natural quadrupeds. Future work will focus on the hardware realization of S-Cheetah for sim-to-real transfer, followed by the integration of imitation learning with canine motion capture (mocap) data to further investigate natural and efficient spinal motion patterns.

## REFERENCES

- [1] R. M. Alexander, *Principles of animal locomotion*. Princeton university press, 2003.
- [2] M. Hildebrand, “The quadrupedal gaits of vertebrates,” *Bioscience*, vol. 39, no. 11, p. 766, 1989.
- [3] T. Kamimura, S. Aoi, Y. Higurashi, N. Wada, K. Tsuchiya, and F. Matsuno, “Dynamical determinants enabling two different types of flight in cheetah gallop to enhance speed through spine movement,” *Scientific Reports*, vol. 11, no. 1, p. 9631, 2021.
- [4] B. Katz, J. D. Carlo, and S. Kim, “Mini cheetah: A platform for pushing the limits of dynamic quadruped control,” in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 6295–6301.
- [5] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter, “Learning agile and dynamic motor skills for legged robots,” *Science Robotics*, vol. 4, no. 26, p. eaa5872, 2019. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.aau5872>
- [6] D. Hoeller, N. Rudin, D. Sako, and M. Hutter, “Anymal parkour: Learning agile navigation for quadrupedal robots,” *Science Robotics*, vol. 9, no. 88, p. eadi7566, 2024. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.adi7566>
- [7] H. Kim, H. Oh, J. Park, Y. Kim, D. Youm, M. Jung, M. Lee, and J. Hwangbo, “High-speed control and navigation for quadrupedal robots on complex and discrete terrain,” *Science Robotics*, vol. 10, no. 102, p. eads6192, 2025. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.ads6192>
- [8] Xinhua News Agency. (2025, jul) Chinese-made robotic dog hits 10.3 m/s, rivals human sprinters. Xinhua. [Online; accessed 9-January-2026]. [Online]. Available: <http://english.news.cn/20250707/cc729a6a50854206ae621f6e142886d5c.html>
- [9] Guinness World Records. (2025, may) Fastest 100 m by a quadrupedal robot. Guinness World Records. [Online; accessed 9-January-2026]. [Online]. Available: <https://www.guinnessworldrecords.com/world-records/625586-fastest-100-m-by-a-quadrupedal-robot>
- [10] S. Bhattacharya, A. Singla, Abhimanyu, D. Dholakiya, S. Bhatnagar, B. Amrutur, A. Ghosal, and S. Kolathaya, “Learning active spine behaviors for dynamic and efficient locomotion in quadruped robots,” in *2019 28th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 2019, pp. 1–6.
- [11] J. Caporale, Z. Feng, S. Rozen-Levy, A. M. Carter, and D. E. Koditschek, “Twisting spine or rigid torso: Exploring quadrupedal morphology via trajectory optimization,” in *2023 IEEE International Conference on Robotics and Automation*, no. 886, 2023.
- [12] Q. Qian, X. Wei, Z. Zhang, J. Tu, Y. Zhang, T. Hou, X. Gao, P. Zhai, and L. Zhang, “A robust quadruped robot with twisting waist for flexible motions,” *arXiv preprint arXiv:2410.05884*, 2024.
- [13] K. Yoneda, K. Kawaharazuka, T. Suzuki, T. Hattori, and K. Okada, “Kleiyn : A quadruped robot with an active waist for both locomotion and wall climbing,” in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025, pp. 8783–8789.
- [14] Z. Yang, Z. Xu, H. Rong, S. Mo, Y. Chen, Z. Chen, T. Wang, and H. Cheng, “Enhancing the flexibility of a quadruped robot with a 2-dof active spine using nonlinear model predictive control,” in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025, pp. 12 526–12 532.
- [15] J. Wang, J. Cheng, J. Hu, W. Gao, and S. Zhang, “Spined torso renders advanced mobility for quadrupedal locomotion,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 1457–1463.
- [16] Y. Wu, Z. Zhou, R. Liu, L. Zheng, and H. Cheng, “A Generic Continuous Multi-Joint Spinal Robotic System for Agile and Accurate Behaviors with GNN-MPC method,” in *Proceedings of Robotics: Science and Systems*, LosAngeles, CA, USA, June 2025.
- [17] J. Hu, W. Gao, X. Zhang, J. Cheng, and S. Zhang, “Predictive control of a spined quadrupedal robot based on a dual rigid-body model,” in *2024 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, 2024, pp. 292–298.
- [18] W. Li, Z. Zhou, and H. Cheng, “Dynamic locomotion of a quadruped robot with active spine via model predictive control,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 1185–1191.
- [19] R. I. Belyaev, P. Nikolskaia, A. V. Bushuev, A. A. Panyutina, D. A. Kozhanova, and N. E. Prilepskaya, “Running, jumping, hunting, and scavenging: Functional analysis of vertebral mobility and backbone properties in carnivorans,” *Journal of Anatomy*, vol. 244, no. 2, pp. 205–231, 2024.
- [20] T. R. Kane and M. Scher, “A dynamical explanation of the falling cat phenomenon,” *International journal of solids and structures*, vol. 5, no. 7, pp. 663–670, 1969.

- [21] C. Schwarke, M. Mittal, N. Rudin, D. Hoeller, and M. Hutter, "Rsl-rl: A learning library for robotics research," *arXiv preprint arXiv:2509.10771*, 2025.
- [22] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [23] Y. G. Alqaham, J. Cheng, and Z. Gan, "16 ways to gallop: Energetics and body dynamics of high-speed quadrupedal gaits," in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025, pp. 16 168–16 174.
- [24] Z. Fan, "robot\_lab: RL extension library for robots, based on isaacsim." 2024. [Online]. Available: [https://github.com/fan-ziqi/robot\\_lab](https://github.com/fan-ziqi/robot_lab)