

Bipedalism for Quadrupedal Robots: Versatile Loco-Manipulation through Risk-Adaptive Reinforcement Learning

Zhang, Yuyou; Corcodel, Radu

TR2025-140 October 01, 2025

Abstract

Loco-manipulation of quadrupedal robots has broadened robotic applications, but using legs as manipulators often compromises locomotion, while mounting arms complicates the system. To mitigate this issue, we introduce bipedalism for quadrupedal robots, thus freeing the front legs for versatile interactions with the environment. We propose a risk-adaptive distributional Reinforcement Learning (RL) framework designed for quadrupedal robots walking on their hind legs, balancing worst-case conservativeness with optimal performance in this inherently unstable task. During training, the adaptive risk preference is dynamically adjusted based on the uncertainty of the return, measured by the coefficient of variation of the estimated return distribution. Extensive experiments in simulation show our method’s superior performance over baselines. Real-world deployment on a Unitree Go2 robot further demonstrates the versatility of our policy, enabling tasks like cart pushing, obstacle probing, and payload transport, while showcasing robustness against challenging dynamics and external disturbances.

IEEE-RAS International Conference on Humanoid Robots 2025

Bipedalism for Quadrupedal Robots: Versatile Loco-Manipulation through Risk-Adaptive Reinforcement Learning

Yuyou Zhang¹, Radu Corcodel², Ding Zhao¹

Abstract—Loco-manipulation of quadrupedal robots has broadened robotic applications, but using legs as manipulators often compromises locomotion, while mounting arms complicates the system. To mitigate this issue, we introduce bipedalism for quadrupedal robots, thus freeing the front legs for versatile interactions with the environment. We propose a risk-adaptive distributional Reinforcement Learning (RL) framework designed for quadrupedal robots walking on their hind legs, balancing worst-case conservativeness with optimal performance in this inherently unstable task. During training, the adaptive risk preference is dynamically adjusted based on the uncertainty of the return, measured by the coefficient of variation of the estimated return distribution. Extensive experiments in simulation show our method’s superior performance over baselines. Real-world deployment on a Unitree Go2 robot further demonstrates the versatility of our policy, enabling tasks like cart pushing, obstacle probing, and payload transport, while showcasing robustness against challenging dynamics and external disturbances.

I. INTRODUCTION

In recent years, the field of quadrupedal robots has made remarkable progress. In terms of locomotion, improved capabilities of traversing various terrains and outdoor environments were developed [1]–[11]. Manipulation skills [12]–[16] and specialized abilities, such as ball shooting, dribbling, catching, and goalkeeping [17]–[20], have been further studied to expand the real-world applicability of quadrupedal robots. These components enable legged robots flexible operations in unstructured environments, and are the building blocks of general sensorimotor skills that facilitate meaningful interactions between legged robots and their surroundings.

One way to enable interactions with the environment is to decouple the locomotion and manipulation components by equipping the quadrupedal robots with top-mounted robotic arms [12], [21], [22] or claws [23]. However, adding robotic arms to quadrupedal robots significantly limits their applicability due to increased weight, energy demand, and additional spatial constraints. Inspired by bipedalism in human evolution [24], we adopt *bipedal gait* for quadrupedal robots to free up their front legs, which are typically used for locomotion in a quadrupedal gait, and repurpose them for manipulation tasks, such as pushing, environment probing,

payload carrying and other tasks that require non-prehensile interactions with the surrounding environment.

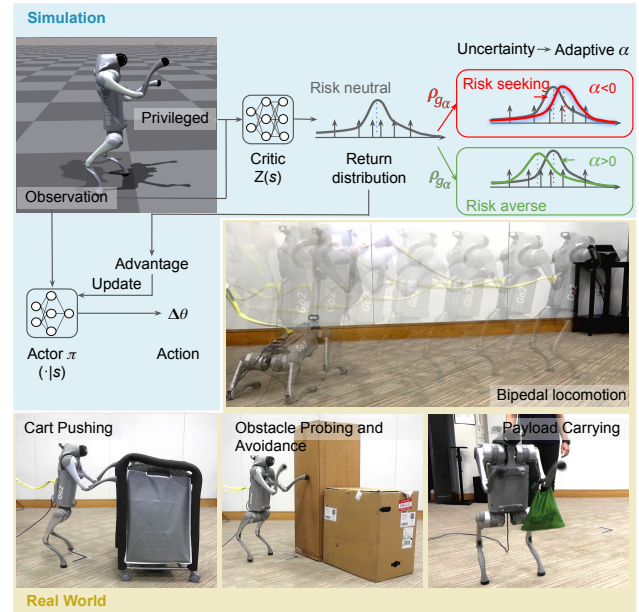


Fig. 1. Risk-adaptive distributional RL framework overview for bipedal locomotion on a Unitree Go2 robot. Using the distortion risk measures ρ_{g_α} , the policy tends to be optimistic (red) in the well-explored states when the uncertainty of the return distribution is low, and vice versa (green). Versatile real-world applications such as cart pushing, obstacle probing, and payload carrying, are enabled by a single locomotion policy. Demonstrations are available in our supplementary video.

Bipedal locomotion differs from quadrupedal locomotion in its inherent instability [25] due to the narrower base of support. These challenges can be exacerbated in the real world when dynamics and disturbances are unknown [26]. RL is commonly used to learn a control policy with complex dynamics. RL in standard form is usually risk-neutral and maximizes the expected accumulated return [27], [28]. The sim-to-real gap and unexpected perturbations during deployment can destabilize a risk-neutral policy, which focuses only on expected return, especially when worst-case returns are underrepresented in the return distribution [26]. It is thus crucial to adopt a risk-aware approach to consider these worst-case scenarios, particularly when the unpredictability of real-world deployment poses significant risks.

In this work, we propose a risk-adaptive distributional RL framework, as shown in Figure 1, to learn a robust policy for inherently unstable bipedal locomotion. Specifically, during training, we adapt the risk preference dynamically based on the uncertainty of the return, measured by the coefficient of variation of the estimated return distribution, instead of pre-

*This work was fully supported by Mitsubishi Electric Research Labs (MERL)

¹Department of Mechanical Engineering, Carnegie Mellon University, Pittsburgh, PA 15213 USA, {yuyou, dingzhao}@andrew.cmu.edu

²Mitsubishi Electric Research Labs (MERL), Cambridge, MA 02139 USA, corcodel@merl.com

specifying the risk level for policy learning. Extensive simulation experiments demonstrate the superior performance of our method compared to baseline approaches. In real-world deployment, we showcase loco-manipulation including cart pushing, contact-aware obstacle probing, and payload carrying, highlighting the versatility and robustness of the bipedal locomotion policy. In summary, the main contributions of this work are:

- We introduce a risk-adaptive RL framework for the robust bipedal locomotion of quadrupedal robots.
- We propose a novel uncertainty metric based on the return distribution to adaptively choose the risk level.
- We demonstrate robust real-world applications with bipedal locomotion under external force, and highlight three representative tasks including cart pushing, contact-aware obstacle probing, and payload carrying.

II. RELATED WORK

A. Quadrupedal robot locomotion and manipulation

Previous work either equips the legged robot with a mounted robotic arm [12]–[15], [21], [22], [29]–[35] or use one leg as manipulator [17], [18], [36]–[38]. With mounted arms, mobile manipulation requires coordination between the robot arm and the legged robot. Instead of decoupling the manipulation and locomotion controllers like in [29], [34], more recent work seeks to build manipulation-centric whole-body controllers to allow better coordination: [21], [30] formulate a unified whole-body MPC framework, [22] uses RL to train whole-body control policy for end-effector tracking, [31] trains RL policy with vision input, [12] uses diffusion policy to learn from human demonstration, [33] learns a whole-body force control policy to enable compliance and force application. However, adding a mounted arm to a quadrupedal robot increases the load and energy requirements, and unnecessarily adds system complexity.

Loco-manipulation repurposes the robot’s legs for manipulation without changing its embodiment. However, increased manipulation ability comes at the cost of compromised locomotion ability since the robot’s legs are primarily designed for locomotion. For example, most loco-manipulations have the robots stand still on three legs and use one leg as the manipulator [23], [36], [38]–[41]. Inspired by how bipedalism played an important role in human evolution by freeing the hands for manipulation [24], we adopt a bipedal gait to enable flexible loco-manipulation such as bimanual pushing, contact-aware obstacle probing, and payload carrying. Compared to previous work on bipedal locomotion of quadrupedal robots [7], [42], [43], our work further explores meaningful interactions between quadrupedal robots and their surrounding environment, enabled by our robust bipedal locomotion policy.

B. Risk-aware RL for Robot

Risk awareness is essential for the successful real-world deployment of autonomous robots, such as drones [44] and quadrupedal robots [45]–[48]. Distributional RL [49]–[51] models the distribution of returns explicitly, rather than

estimating the value function as the expected return, making it widely applicable in risk-sensitive RL [52], [53].

Several works have used distributional RL to improve the robustness of quadrupedal robot real-world performance. Li et al. [46] introduce a distribution ensemble actor-critic approach and demonstrate improved performance in domain randomization settings. However, the approach is not risk-aware. Shi et al. [47] takes a risk-adaptive perspective but can only switch between a risk-neutral (CVaR₁) policy and a risk-averse (CVaR_{0.5}) policy, resembling switch-mode control. Also, risk-averse learning [47] can ignore high-return strategies [52]. Schneider et al. [45] propose a risk-aware Distributional PPO and demonstrate risk-aware locomotion behavior conditioned on a manually specified risk level at deployment, which relies on human prior knowledge. Different from previous work, our method enables risk adaptiveness during training, allowing the actor policy to internalize risk preference selection automatically. During training, the value estimation becomes conservative in high-uncertainty situations and autonomously shifts toward optimism when the return distribution exhibits low variance. This results in a policy that is both robust and high-performing, without requiring manual risk tuning at deployment.

III. PRELIMINARY

Partially Observable Markov Decision Process. We formulate bipedal locomotion learning as a Partially Observable Markov Decision Process (POMDP) defined by $(S, \mathcal{A}, \mathcal{T}, R, \Omega, O, \gamma)$, where S represents the state space, \mathcal{A} the action space, $\mathcal{T} : S \times \mathcal{A} \mapsto S$ is the transition function, $R : S \times \mathcal{A} \mapsto \mathbb{R}$ is the reward function, Ω is the set of observations, O is the observation function, γ is the discount factor. The objective is to train a policy π^* which maximizes the discounted cumulative reward $\pi^* = \arg \max_{\pi} \mathbb{E}_{s_0 \sim \rho_0, a_t \sim \pi(\cdot|s_t)} \left[\sum_{t \geq 0} \gamma^t r(s_t, a_t) \right]$.

Distributional RL. Distributional RL [49]–[51] learns the value distribution, instead of the expected return as a value function. With policy π , the return is a random variable Z^{π} that represents the cumulated discounted rewards along one trajectory, $Z^{\pi} = \sum_{t=0}^{\infty} \gamma^t R_t$. The value function for many standard RL algorithms is, $V^{\pi}(x) = \mathbb{E}[Z^{\pi}(x)]$ while distributional RL explicitly parameterizes the return distribution with quantile functions [50], [51] or discrete distribution [49]. We adopt a similar parameterization as in QR-DQN [50], where the return distribution is approximated by estimating the quantiles $\tau_1, \dots, \tau_N, \tau_i = i/N, i = 1, \dots, N$, with the parametric model θ ,

$$Z_{\theta}(x) := \frac{1}{N} \sum_{i=1}^N \delta_{\theta_i(x)}, \quad (1)$$

where $\theta_i(x)$ is the i^{th} quantile of the return distribution $Z^{\pi}(x)$, and $\delta_{\theta_i(x)}$ denotes a Dirac function at $\theta_i(x)$.

IV. METHOD

We first introduce the problem setting of the bipedal locomotion task in section IV-A. Due to the robustness

requirement of balance in bipedal locomotion, we adopt distributional RL shown in section IV-B. In section IV-C, we incorporate adaptive risk measures to balance safety in worst-case scenarios with optimal task performance. The framework overview of our proposed method is shown in Figure 1.

A. Control Policy with RL

Observation and Action Space. The bipedal locomotion policy receives observations which include the proprioceptive information, locomotion command, and the last action $a_{t-1} \in \mathbb{R}^{12}$. Proprioceptive information includes joint position $\theta_t \in \mathbb{R}^{12}$ and joint velocity $\dot{\theta}_t \in \mathbb{R}^{12}$ provided by the joint encoders and projected gravity in the robot frame $g_t \in \mathbb{R}^3$ from the IMU. The command $\mathbf{c}_t = [v_x^c, v_y^c, \omega_{yaw}^c, z^c, f^c]$ includes the velocity command specifying the linear velocities in the longitudinal and lateral directions, angular velocity around the vertical axis, base height, and stepping frequency. Privileged observation for the critic network at the training stage includes extra information only available in simulation such as joint friction and restitution coefficient. The dimension of the action space \mathcal{A} is 12, which equals the number of actuators. The predictions of the policy, $\Delta\theta_t \in \mathbb{R}^{12}$, are the joint angles relative to the nominal quadrupedal standing position.

Reward Functions. Reward functions consist of task-specific rewards for bipedal locomotion and auxiliary rewards adapted from [54] to optimize foot contact, action smoothness, energy consumption, joint position, etc. Task-specific reward functions are listed in Table I. In addition to *Base Height* which encourages maintaining base height z^c and *Base Pitch* to promote an upright position, we also use *Upright Balance* to penalize velocity along the z -axis and changes in pitch angle \dot{p} when the robot is upright. Velocity tracking rewards include *Linear Tracking* and *Angular Tracking*, where σ and σ_{yaw} are scaling factors.

Apart from direct tracking reward, we design a *Support Polygon* reward to track the relative position between the base center of mass (CoM) and the support polygon. When the CoM moves ahead of the support polygon, the robot accelerates. Conversely, when the CoM lags behind, the robot decelerates. This mechanism enables continuous balance control while tracking the desired velocity, as shown in Fig. 2. We characterize the relative position between the robot CoM and support polygon by $\arctan(\Delta x_b / \Delta z_b)$. Δx_b and Δz_b are relative positions of the average of two rear feet in the body frame along the x -axis and z -axis. $\arctan(\Delta x_b / \Delta z_b)$ is expected to be positive when the robot needs to accelerate and negative when decelerating. This angle is essentially different than the pitch angle and only degenerates to the pitch angle when the model is simplified to a single inverted pendulum. The total reward is a positive linear function of the task reward, $r_+ \times e^{c \times r_-}$, where r_+ is the sum of positive reward terms and r_- is the sum of negative reward terms, c is coefficient set to 0.02.




Fig. 2. The robot accelerates (a), stays neutral (b), and decelerates (c) by shifting its center of mass (CoM) ahead of, aligned with, or behind the support polygon.

TABLE I
TASK REWARD FUNCTIONS

<i>Base Height</i>	$-(z - z^c)^2$
<i>Base Pitch</i>	$-\cos(p^c - p)$
<i>Upright Balance</i>	$\exp(-v_z^2/\sigma) + \exp(-\dot{p}^2/\sigma_{yaw})$ if is upright, else 0
<i>Linear Tracking</i>	$\exp(- v_{xy} - v_{xy}^c ^2/\sigma)$ if is upright, else 0
<i>Angular Tracking</i>	$\exp(- w_{yaw} - w_{yaw}^c ^2/\sigma_{yaw})$ if is upright, else 0
<i>Support Polygon</i>	$- v_x^c ^2 \left(\frac{\pi}{2} - \arctan(\frac{\Delta x_b}{\Delta z_b}) \right)^2$ if $\arctan(\frac{\Delta x_b}{\Delta z_b}) v_x^c < 0$

B. Risk-aware Distributional PPO

In Distributional PPO, the actor network remains the same as in standard PPO [27] and is trained to optimize the clipped objective,

$$\mathcal{L}(\phi) = \mathbb{E}_t \left[\min(\eta_t(\phi) \hat{A}_t, \text{clip}(\eta_t(\phi), 1 - \epsilon, 1 + \epsilon) \hat{A}_t) \right], \quad (2)$$

where $\eta_t(\phi) = \frac{\pi_\phi(a_t|o_t)}{\pi_{\phi_{\text{old}}}(a_t|o_t)}$, is the probability ratio between the new policy π_ϕ , and old policy $\pi_{\phi_{\text{old}}}$, \hat{A}_t is the estimated advantage computed using the Generalizable Advantage Estimation (GAE) [55]. The critic network differs from standard PPO by predicting the quantiles of the value distribution as in Equation 1, instead of just the expected value. $\theta_i(x) = F_Z^{-1}(\tau_i)$ for $\tau_i = i/N, i = 1, \dots, N$ are N -quantiles of the return distribution $Z^\pi(x)$, and F_Z^{-1} denotes the inverse CDF of the return distribution $Z^\pi(x)$. For a critic network that approximate the return distribution by predicting $\theta_1(x), \dots, \theta_N(x)$, the objective is to minimize the following quantile loss, which effectively minimizes the 1-Wasserstein distance between the empirical distribution $\hat{Z}^\pi(x)$ and the parameterized quantile distribution $Z_\theta^\pi(x)$:

$$\mathcal{L}(\theta) = \mathbb{E}_t \left[\frac{1}{N} \sum_{i=1}^N (\tau - \mathbb{1}_{z_t - \theta_t^i < 0})(z_t - \theta_t^i) \right], \quad (3)$$

where $z_i \sim Z^\pi(x_t)$, and $\theta_t^i = \theta_i(x_t)$.

The value function derived from the return distribution is used to estimate the advantage \hat{A}_t for the actor-network update. Given the return distribution predicted by the critic network, the risk-neutral value function is calculated as,

$$V(x) = \sum_{i=1}^N (\tau_i - \tau_{i-1}) \theta_i(x) = \sum_{i=1}^N \frac{1}{N} \theta_i(x). \quad (4)$$

To learn a risk-aware policy, the distortion risk measures ρ_{g_α} associated with the distortion function g_α is applied to

the return distribution. Conditional Value at Risk (CVaR) is commonly applied to have risk-averse behaviors since only the left tail is quantified. We apply Wang’s metric [56] so that the risk preference can be adjusted from averse ($\alpha > 0$) to seeking ($\alpha < 0$) as in equation 5,

$$g_{\alpha}^{\text{Wang}}(\tau) = \Phi(\Phi^{-1}(\tau) + \alpha). \quad (5)$$

Then the new value function is given by the distorted return distribution as in Equation 6.

$$V_{\alpha}(x) = \rho_{g_{\alpha}}(Z_{\theta}^{\pi}(x)) = \sum_{i=1}^N (g_{\alpha}(\tau_i) - g_{\alpha}(\tau_{i-1})) \theta_i(x). \quad (6)$$

When $\alpha > 0$, the calculation of the value function is conservative, by assigning more weight to worst-case left tails, while $\alpha < 0$ makes the calculation optimistic by assigning more weight to higher returns.

C. Uncertainty Modeling and Adaptive Risk Level

We assume the transition is deterministic in this POMDP, then the aleatory uncertainty comes from partial observation of the state and domain randomization, while the epistemic uncertainty arises from the lack of environment knowledge to make informed predictions. Optimism in the face of epistemic uncertainty encourages exploration, allowing for the collection of more informative data to maximize long-term returns [28]. However, conservativeness is essential to ensure worst-case performance when aleatory uncertainty is present. We model the uncertainty as the uncertainty of the parameterized distribution $Z_{\theta}^{\pi}(x)$ predicted by the critic network, represented by the Coefficient of Variation (CV),

$$CV_{Z_{\theta}^{\pi}(x)} = \frac{\sqrt{\text{Var}(Z_{\theta}^{\pi}(x))}}{\mathbb{E}Z_{\theta}^{\pi}(x)} = \frac{\sigma}{\mu}. \quad (7)$$

This normalized measure allows for the comparison of dispersion across different return distributions, even if the means are drastically different from each other. It is particularly advantageous over non-normalized measures, such as right truncated variance (RTV) used in [44] and interquartile range (IQR) used in [47], as the mean of the return distribution tends to increase significantly during training due to increased rewards.

We formulate the parameter α of the distortion function $g_{\alpha}^{\text{Wang}}(\tau)$ as a function of the modeled uncertainty $CV_{Z_{\theta}^{\pi}(x)}$ and training steps t ,

$$\alpha_t = (\alpha_0 - \alpha_T) e^{-\frac{t/T}{CV_t}} + \alpha_T, \quad (8)$$

where T is the total training steps. And CV_t is the average over batch of $CV_{Z_{\theta}^{\pi}(x)}$ at step t . With $\alpha_0 - \alpha_T > 0$, the policy begins conservatively and becomes increasingly optimistic as training progresses. The dependence of α_t on CV_t directs the policy to be more conservative when uncertainty is high. The time-dependent coefficient $\frac{t}{T}$ allows α_t to be progressively more influenced by the uncertainty CV_t as training advances.

TABLE II

PROPOSED METHOD AND BASELINES SETTINGS			
Method	Distributional	Fixed Risk	Adaptive Risk
<i>DPPO_adaptive</i> (ours)	✓	✗	✓
<i>DPPO_neutral</i> [45]	✓	$\alpha = 0$	✗
<i>DPPO_averse</i> [45]	✓	$\alpha > 0$	✗
<i>DPPO_seeking</i> [45]	✓	$\alpha < 0$	✗
<i>PPO</i> [27]	✗	✗	✗

V. EXPERIMENTS

We evaluate our method across several experiments. In Section V-A, we demonstrate our method has higher training rewards compared to baseline methods and ablations. In Section V-B, we evaluate velocity tracking error across in-distribution and out-of-distribution target velocities, as well as under external forces. In Section V-C, we analyze uncertainty and risk modeling to highlight the importance of risk-adaptive learning. Finally, in Section V-D, we deploy our bipedal locomotion policy on the Unitree Go2 robot, showing that a single policy enables robust bipedal locomotion and versatile loco-manipulation capabilities.

Simulation setup We use Isaac Gym [57] to train the bipedal locomotion policy based on the open-source framework in [54]. The target velocity range is $[-0.8, 0.8]$ m/s for v_x^c and is $[-0.4, 0.4]$ m/s for v_y^c , and $[-1, 1]$ rad for ω_{yaw} . The critic network and actor network both have hidden dimensions [512, 256, 128]. The output layer size of the critic network is 64, to predict $N = 64$ quantiles of the return distribution.

We initialize our risk-adaptive DPPO with neutral initial risk $\alpha_0 = 0$ to maintain stability and avoid early catastrophic failures during training. However, remaining too conservative may hinder exploration and limit performance improvements. Therefore, we gradually shift to a more optimistic risk preference $\alpha_T = -0.2$, encouraging the agent to explore high-return strategies and accelerate learning. We select these parameters to avoid over-conservativeness early on, which may prevent discovering successful quadrupedal-to-bipedal transitions, and to limit excessive optimism later, which could destabilize training. We train 4000 agents in parallel for 20k iterations on an NVIDIA RTX 4090 GPU, which takes approximately 5 hours. We compare our method to baselines in Table II. All methods share the same hyperparameters if applicable.

Hardware setup We use the Unitree Go2 robot for real-world experiments. The computations are performed on a host computer. The policy runs at 50Hz and the robot receives the joint position command from the host computer. Target joint angles were tracked using a PD controller with gains set to $K_p = 25$ and $K_d = 0.6$.

A. Training Performance

Baseline comparison Figure 3 shows the comparison between proposed risk-adaptive distributional PPO (DPPO) and baselines in Table II. Parameter α in the distortion function $g_{\alpha}^{\text{Wang}}(\tau)$ for *DPPO_averse* is 0.2, and for *DPPO_seeking* is -0.2 , which equals to the final risk α_T for *DPPO_adaptive*.




Fig. 3. Learning curves of proposed method (*DPPO_adaptive*) against baselines listed in Table II. The rewards are averaged over three seeds, and the shaded region represents the standard error.

Our method consistently outperforms the baselines in both velocity tracking and total reward. Risk-neutral DPPO and PPO perform similarly and both achieve a lower total reward compared to DPPO with adaptive risk. Risk-seeking DPPO fails to learn a locomotion policy, resulting in near-zero rewards after 20k steps, and diverging policy entropy, indicating that constant risk-seeking can lead to catastrophic failures. In contrast, risk-adaptive DPPO exhibits higher policy entropy during training compared to the risk-neutral baselines (Risk-neutral DPPO and PPO), suggesting that our method encourages exploration of diverse actions, while the risk-neutral policies are less exploratory. Risk-averse DPPO achieves the lowest velocity tracking reward but ranks second-to-last in total reward, due to its risk-averse strategy, which minimizes the accumulation of negative penalties that contribute to the total reward.

Ablations of reward functions To assess the impact of proposed reward functions in Table I, we compare our method to *DPPO_adaptive_w/o_support* and *DPPO_adaptive_w/o_balance*, where, in each case, one of the task reward functions is removed. As shown in Figure 4, the absence of the *Support Polygon* reward function leads to a significant drop in linear velocity tracking performance, and the *Upright Balance* reward function enhances overall bipedal locomotion performance.




Fig. 4. Learning curves of our method and reward function ablations. The rewards are averaged over three seeds, and the shaded region represents the standard error.

B. Tracking Error Evaluation

We evaluate the learned policy based on success rate and velocity tracking error. The evaluation is averaged across 4000 environments, each with an episode length of 1000 steps. An episode is considered successful if it does not terminate early due to the robot crashing. The tracking error is calculated as the Root mean square error (RMSE) across all evaluation environments and episode steps.

Across Varying Target Velocities With the training target velocity v_x sampled in the range of $[-0.8, 0.8]$ m/s, we




Fig. 5. Success Rate (a) and X Tracking Error (b) across target velocities ranging from -1.0 m/s to 1.0 m/s. Comparison of *DPPO_adaptive* with three baseline methods.

use in-distribution velocities of $[\pm 0.8, \pm 0.5, \pm 0.2, 0.0]$ m/s and out-of-distribution (OOD) velocities of ± 1.0 m/s as the evaluation target velocities. We show that risk-adaptive DPPO outperforms baselines with the highest success rate and lowest tracking error in Table III. More specifically, risk-adaptive DPPO achieves the lowest tracking error for target velocities of -0.25 m/s or higher in Figure 5. For target velocities below -0.25 m/s, PPO and risk-averse DPPO perform better, suggesting that backward velocity tracking may require a more conservative policy. This is consistently indicated by the generally lower success rate for backward tracking compared to forward velocity tracking. Despite this, we did not fully explore the potential of our method by tuning the initial and final risk levels, as forward velocity tracking is more common in real-world deployments.

TABLE III
AVERAGE VELOCITY TRACKING ERROR AND SUCCESS RATE ACROSS DIFFERENT TARGET VELOCITY

Method	Success Rate \uparrow	x_RMSE \downarrow	y_RMSE \downarrow
DPPO_adaptive (ours)	0.964 \pm 0.0044	0.128 \pm 0.0100	0.072 \pm 0.0002
DPPO_averse [45]	0.922 \pm 0.0142	0.135 \pm 0.0088	0.080 \pm 0.0028
DPPO_neutral [45]	0.962 \pm 0.0006	0.139 \pm 0.0035	0.074 \pm 0.0010
PPO	0.962 \pm 0.0011	0.149 \pm 0.0003	0.077 \pm 0.0002

Even though backward target velocity prefers a more conservative policy, we show that risk-adaptive DPPO with a risk-seeking tendency outperforms neutral and risk-averse baselines, showing significant generalizability when evaluated with OOD velocity command ± 1.0 m/s, including negative backward velocity -1 m/s, as shown in Table IV.

Under external force We evaluate the performance under external force to further assess the robustness of our proposed method, as shown in Table V. A 10N external force was applied downward on each of the robot's forearms with a velocity command of 1m/s. Our method achieves the highest success rate, nearly doubling the second-best, and also exhibits the smallest drop in success rate compared to conditions without external force. Although risk-averse DPPO shows a slightly smaller X tracking error, this doesn't indicate better performance, as its success rate is only half that of our method. The lower error is likely induced by early-terminated episodes, which could have exhibited significantly higher tracking errors if they had not failed.

TABLE IV
OUT OF DISTRIBUTION TARGET VELOCITY TRACKING ERROR

Method	1m/s			-1m/s		
	Success Rate \uparrow	x_RMSE \downarrow	y_RMSE \downarrow	Success Rate \uparrow	x_RMSE \downarrow	y_RMSE \downarrow
DPPO_adaptive	0.976 \pm 0.009	0.151 \pm 0.007	0.072 \pm 0.001	0.880 \pm 0.009	0.194 \pm 0.039	0.087 \pm 0.007
DPPO_averse [45]	0.911 \pm 0.088	0.175 \pm 0.021	0.086 \pm 0.010	0.819 \pm 0.010	0.183 \pm 0.026	0.094 \pm 0.009
DPPO_neutral [45]	0.956 \pm 0.018	0.154 \pm 0.014	0.083 \pm 0.005	0.875 \pm 0.003	0.235 \pm 0.014	0.088 \pm 0.008
PPO	0.929 \pm 0.004	0.174 \pm 0.005	0.097 \pm 0.004	0.864 \pm 0.006	0.249 \pm 0.005	0.098 \pm 0.001

TABLE V
VELOCITY TRACKING ERROR UNDER EXTERNAL FORCE

Method	Success Rate \uparrow	Success Rate Drop \downarrow	x_RMSE \downarrow	y_RMSE \downarrow
DPPO_adaptive	0.601 \pm 0.418	38.42%	0.272 \pm 0.063	0.097 \pm 0.021
DPPO_averse [45]	0.327 \pm 0.268	64.11%	0.247 \pm 0.025	0.120 \pm 0.012
DPPO_neutral [45]	0.016 \pm 0.009	98.32%	0.337 \pm 0.026	0.144 \pm 0.004
PPO	0.029 \pm 0.003	89.77%	0.300 \pm 0.001	0.118 \pm 0.004

Our method enables robust adaptation to disturbances while maintaining high performance by dynamically adjusting risk based on value function uncertainty during training, which is further explained in Section V-C.

C. Distribution Uncertainty Visualization

To study how risk adaptiveness improves performance, we plot the uncertainty of the estimated return distribution, denoted by the coefficient of variation (CV) in Figure 7. The robot attempts to follow the 0m/s command, where external forces are applied for 0.5s at 2s intervals, with magnitudes ranging from 20N to 100N. Each force exertion corresponds to an increase in both velocity deviation and uncertainty, showing temporary instability. After each peak, the uncertainty gradually declines, reflecting the policy effectively learns to regain stability. With the coefficient of variation as a metric of the uncertainty of critic network distribution, we validate that the model identifies these higher-risk or uncertain situations caused by external perturbations.

At $t = 0$, the uncertainty is high because the quadrupedal-to-bipedal transition involves less frequently visited states, resulting in greater uncertainty in the critic’s value estimation. More optimistic actions can be taken in well-explored states, such as bipedal tracking without disturbances. This underscores the importance of risk-adaptive DPPO in balancing conservatism in high-uncertainty states with optimism in well-explored states.

D. Real World Deployment

We deploy our policy on the Go2 robot in the real world, showcasing a single policy that enables versatile locomotion capabilities, as illustrated in Figure 6. This policy enables not only basic locomotion such as forward, backward, and turning maneuvers but also supports complex interactions to further demonstrate its versatility. The robot can effectively fulfill tasks such as cart pushing, obstacle probing, and payload carrying. All of the real-world tasks introduce external forces that could destabilize the robot and require the robustness of the policy. Pushing a cart demands a robust loco-manipulation policy that can adjust the force applied and stabilize the body accordingly. Obstacle probing requires the robot to recover from an unstable state when it runs into an obstacle and probes the obstacle with its front legs. Carrying a payload increases the weight and shifts the

robot’s center of mass, necessitating dynamic balance and stability. Remarkably, the success of these real-world tasks is a direct result of the single bipedal locomotion policy, without requiring extensive task-specific training. This showcases the robustness and versatility of our bipedal locomotion policy and validates the effectiveness of our proposed risk-adaptive learning framework.

VI. CONCLUSION

In this work, we introduce a risk-adaptive distributional RL framework for quadrupedal robots, enabling robust bipedal locomotion and versatile interactions with complex environments. Through extensive simulation and real-world experiments, we validate the robustness and adaptability of this framework, which is grounded in modeling the return distribution for risk-adaptive learning. Refining the adaptation strategy for more flexible risk management could further enhance its performance. Future work may also focus on high-level planning and closed-loop control to facilitate long-horizon tasks.

REFERENCES

- [1] N. Rudin, D. Hoeller, P. Reist, and M. Hutter, “Learning to walk in minutes using massively parallel deep reinforcement learning,” in *Conference on Robot Learning*. PMLR, 2022, pp. 91–100.
- [2] A. Kumar, Z. Fu, D. Pathak, and J. Malik, “Rma: Rapid motor adaptation for legged robots,” *arXiv preprint arXiv:2107.04034*, 2021.
- [3] 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.
- [4] J. Long, Z. Wang, Q. Li, L. Cao, J. Gao, and J. Pang, “Hybrid internal model: Learning agile legged locomotion with simulated robot response,” in *The Twelfth International Conference on Learning Representations*, 2024.
- [5] E. Chane-Sane, J. Amigo, T. Flayols, L. Righetti, and N. Mansard, “Soloparkour: Constrained reinforcement learning for visual locomotion from privileged experience,” in *8th Annual Conference on Robot Learning*, 2024. [Online]. Available: <https://openreview.net/forum?id=DSdAEsEGhE>
- [6] S. Chen, Z. Wan, S. Yan, C. Zhang, W. Zhang, Q. Li, D. Zhang, and F. U. D. Farrukh, “SLR: Learning quadruped locomotion without privileged information,” in *8th Annual Conference on Robot Learning*, 2024. [Online]. Available: <https://openreview.net/forum?id=RMkdcKK7jq>
- [7] J. Long, W. Yu, Q. Li, Z. Wang, D. Lin, and J. Pang, “Learning h-infinity locomotion control,” in *8th Annual Conference on Robot Learning*, 2024. [Online]. Available: <https://openreview.net/forum?id=uMZ2jnZUDX>
- [8] A. L. Mitchell, W. Merkt, A. Papatheodorou, I. Havoutis, and I. Posner, “Gaitor: Learning a unified representation across gaits for real-world quadruped locomotion,” in *8th Annual Conference on Robot Learning*, 2024. [Online]. Available: <https://openreview.net/forum?id=ySI0tBYxpz>
- [9] J. Ren, Y. Liu, Y. Dai, J. Long, and G. Wang, “TOP-nav: Legged navigation integrating terrain, obstacle and proprioception estimation,” in *8th Annual Conference on Robot Learning*, 2024. [Online]. Available: <https://openreview.net/forum?id=O05t1Qt2d5>
- [10] F. Zargarbashi, J. Cheng, D. Kang, R. Sumner, and S. Coros, “Robotkeyframing: Learning locomotion with high-level objectives via mixture of dense and sparse rewards,” in *8th Annual Conference on Robot Learning*, 2024. [Online]. Available: <https://openreview.net/forum?id=wcbRhPnOei>
- [11] R. Yang, Z. Chen, J. Ma, C. Zheng, Y. Chen, Q. Nguyen, and X. Wang, “Generalized animal imitator: Agile locomotion with versatile motion prior,” in *8th Annual Conference on Robot Learning*, 2024. [Online]. Available: <https://openreview.net/forum?id=9XV3dBqfcf>



Fig. 6. Snapshots of bipedal loco-manipulation in the real world. From top to bottom, the images showcase the quadrupedal to bipedal transition, bipedal locomotion, and versatile interactions. Each row is marked with its respective timestamp for chronological analysis. Additional demonstrations, including velocity tracking at various speeds, clockwise and counterclockwise turns, and maintaining a stationary position, can be found in the supplementary video.




Fig. 7. Uncertainties represented by the Coefficient of Variance (CV) visualized during evaluation with a 0 m/s target velocity, external forces are applied for 0.5s at 2s intervals.

- [12] H. Ha, Y. Gao, Z. Fu, J. Tan, and S. Song, "UMI-on-legs: Making manipulation policies mobile with a manipulation-centric whole-body controller," in *8th Annual Conference on Robot Learning*, 2024. [Online]. Available: <https://openreview.net/forum?id=3i7j8ZPnbm>
- [13] J. P. Sleiman, M. Mittal, and M. Hutter, "Guided reinforcement learning for robust multi-contact loco-manipulation," in *8th Annual Conference on Robot Learning*, 2024. [Online]. Available: <https://openreview.net/forum?id=9aZ4ehSTRc>
- [14] R. Mendonca, E. Panov, B. Bucher, J. Wang, and D. Pathak, "Continuously improving mobile manipulation with autonomous real-world RL," in *8th Annual Conference on Robot Learning*, 2024. [Online]. Available: <https://openreview.net/forum?id=46SluHKOe9>
- [15] M. Zhang, Y. Ma, T. Miki, and M. Hutter, "Learning to open and traverse doors with a legged manipulator," in *8th Annual Conference on Robot Learning*, 2024. [Online]. Available: <https://openreview.net/forum?id=VoC3wF6fbh>
- [16] J. Bruedigam, A. A. Abbas, M. Sorokin, K. Fang, B. Hung, M. Guru, S. G. Sosnowski, J. Wang, S. Hirche, and S. L. Cleach, "A versatile planner for learning dexterous and whole-body manipulation," in *8th Annual Conference on Robot Learning*, 2024. [Online]. Available: <https://openreview.net/forum?id=vobaOY0qDI>
- [17] Y. Ji, Z. Li, Y. Sun, X. B. Peng, S. Levine, G. Berseth, and K. Sreenath, "Hierarchical reinforcement learning for precise soccer shooting skills using a quadrupedal robot," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 1479–1486.
- [18] Y. Ji, G. B. Margolis, and P. Agrawal, "Dribblebot: Dynamic legged manipulation in the wild," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 5155–5162.
- [19] B. Forrai, T. Miki, D. Gehrig, M. Hutter, and D. Scaramuzza, "Event-based agile object catching with a quadrupedal robot," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 12 177–12 183.
- [20] X. Huang, Z. Li, Y. Xiang, Y. Ni, Y. Chi, Y. Li, L. Yang, X. B. Peng, and K. Sreenath, "Creating a dynamic quadrupedal robotic goalkeeper with reinforcement learning," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 2715–2722.
- [21] J.-P. Sleiman, F. Farshidian, and M. Hutter, "Versatile multicontact planning and control for legged loco-manipulation," *Science Robotics*, vol. 8, no. 81, p. eadg5014, 2023.
- [22] Z. Fu, X. Cheng, and D. Pathak, "Deep whole-body control: learning a unified policy for manipulation and locomotion," in *Conference on Robot Learning*. PMLR, 2023, pp. 138–149.
- [23] C. Lin, X. Liu, Y. Yang, Y. Niu, W. Yu, T. Zhang, J. Tan, B. Boots, and D. Zhao, "Locoman: Advancing versatile quadrupedal dexterity with lightweight loco-manipulators," *arXiv preprint arXiv:2403.18197*, 2024.
- [24] K. D. Hunt, "The evolution of human bipedality: ecology and functional morphology," *Journal of human evolution*, vol. 26, no. 3, pp. 183–202, 1994.
- [25] J. W. Grizzle, C. Chevallereau, R. W. Sinnet, and A. D. Ames, "Models, feedback control, and open problems of 3d bipedal robotic walking," *Automatica*, vol. 50, no. 8, pp. 1955–1988, 2014.
- [26] P. Akella, A. Dixit, M. Ahmadi, L. Lindemann, M. P. Chapman, G. J. Pappas, A. D. Ames, and J. W. Burdick, "Risk-aware robotics: Tail risk measures in planning, control, and verification," *arXiv preprint arXiv:2403.18972*, 2024.
- [27] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [28] B. O'Donoghue, "Efficient exploration via epistemic-risk-seeking policy optimization," in *Proceedings of the 40th International Conference on Machine Learning*, ser. ICML'23. JMLR.org, 2023.
- [29] C. D. Bellicoso, K. Krämer, M. Stäubli, D. Sako, F. Jenelten, M. Bjelonic, and M. Hutter, "Alma-articulated locomotion and manipulation for a torque-controllable robot," in *2019 International conference on robotics and automation (ICRA)*. IEEE, 2019, pp. 8477–8483.
- [30] J.-P. Sleiman, F. Farshidian, M. V. Minniti, and M. Hutter, "A unified mpc framework for whole-body dynamic locomotion and manipula-

- tion,” *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4688–4695, 2021.
- [31] M. Liu, Z. Chen, X. Cheng, Y. Ji, R.-Z. Qiu, R. Yang, and X. Wang, “Visual whole-body control for legged loco-manipulation,” in *8th Annual Conference on Robot Learning*, 2024. [Online]. Available: <https://openreview.net/forum?id=cT2N3p1AcE>
 - [32] N. Yokoyama, A. Clegg, J. Truong, E. Undersander, T.-Y. Yang, S. Arnaud, S. Ha, D. Batra, and A. Rai, “Asc: Adaptive skill coordination for robotic mobile manipulation,” *IEEE Robotics and Automation Letters*, vol. 9, no. 1, pp. 779–786, 2023.
 - [33] T. Portela, G. B. Margolis, Y. Ji, and P. Agrawal, “Learning force control for legged manipulation,” *arXiv preprint arXiv:2405.01402*, 2024.
 - [34] Y. Ma, F. Farshidian, T. Miki, J. Lee, and M. Hutter, “Combining learning-based locomotion policy with model-based manipulation for legged mobile manipulators,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2377–2384, 2022.
 - [35] J. Zhang, N. Gireesh, J. Wang, X. Fang, C. Xu, W. Chen, L. Dai, and H. Wang, “Gamma: Graspability-aware mobile manipulation policy learning based on online grasping pose fusion,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 1399–1405.
 - [36] Z. He, K. Lei, Y. Ze, K. Sreenath, Z. Li, and H. Xu, “Learning visual quadrupedal loco-manipulation from demonstrations,” *arXiv preprint arXiv:2403.20328*, 2024.
 - [37] X. Huang, Q. Liao, Y. Ni, Z. Li, L. Smith, S. Levine, X. B. Peng, and K. Sreenath, “Hilma-res: A general hierarchical framework via residual rl for combining quadrupedal locomotion and manipulation,” *arXiv preprint arXiv:2407.06584*, 2024.
 - [38] P. Arm, M. Mittal, H. Kolvenbach, and M. Hutter, “Pedipulate: Enabling manipulation skills using a quadruped robot’s leg,” in *41st IEEE Conference on Robotics and Automation (ICRA 2024)*, 2024.
 - [39] X. He, C. Yuan, W. Zhou, R. Yang, D. Held, and X. Wang, “Visual manipulation with legs,” in *8th Annual Conference on Robot Learning*, 2024. [Online]. Available: <https://openreview.net/forum?id=E4K3yLQQ7s>
 - [40] Y. Ouyang, J. Li, Y. Li, Z. Li, C. Yu, K. Sreenath, and Y. Wu, “Long-horizon locomotion and manipulation on a quadrupedal robot with large language models,” *arXiv preprint arXiv:2404.05291*, 2024.
 - [41] X. Cheng, A. Kumar, and D. Pathak, “Legs as manipulator: Pushing quadrupedal agility beyond locomotion,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 5106–5112.
 - [42] Y. Li, J. Li, W. Fu, and Y. Wu, “Learning agile bipedal motions on a quadrupedal robot,” *arXiv preprint arXiv:2311.05818*, 2023.
 - [43] Z. Su, X. Huang, D. Ordoñez-Apaez, Y. Li, Z. Li, Q. Liao, G. Turrissi, M. Pontil, C. Semini, Y. Wu *et al.*, “Leveraging symmetry in rl-based legged locomotion control,” *arXiv preprint arXiv:2403.17320*, 2024.
 - [44] C. Liu, E.-J. van Kampen, and G. C. De Croon, “Adaptive risk-tendency: Nano drone navigation in cluttered environments with distributional reinforcement learning,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 7198–7204.
 - [45] L. Schneider, J. Frey, T. Miki, and M. Hutter, “Learning risk-aware quadrupedal locomotion using distributional reinforcement learning,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 11 451–11 458.
 - [46] S. Li, Y. Pang, P. Bai, J. Li, Z. Liu, S. Hu, L. Wang, and G. Wang, “Learning locomotion for quadruped robots via distributional ensemble actor-critic,” *IEEE Robotics and Automation Letters*, 2024.
 - [47] J. Shi, C. Bai, H. He, L. Han, D. Wang, B. Zhao, M. Zhao, X. Li, and X. Li, “Robust quadrupedal locomotion via risk-averse policy learning,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 11 459–11 466.
 - [48] D. D. Fan, A.-A. Agha-Mohammadi, and E. A. Theodorou, “Learning risk-aware costmaps for traversability in challenging environments,” *IEEE robotics and automation letters*, vol. 7, no. 1, pp. 279–286, 2021.
 - [49] M. G. Bellemare, W. Dabney, and R. Munos, “A distributional perspective on reinforcement learning,” in *International conference on machine learning*. PMLR, 2017, pp. 449–458.
 - [50] W. Dabney, M. Rowland, M. Bellemare, and R. Munos, “Distributional reinforcement learning with quantile regression,” in *Proceedings of the AAAI conference on artificial intelligence*, vol. 32, no. 1, 2018.
 - [51] W. Dabney, G. Ostrovski, D. Silver, and R. Munos, “Implicit quantile networks for distributional reinforcement learning,” in *International conference on machine learning*. PMLR, 2018, pp. 1096–1105.
 - [52] I. Greenberg, Y. Chow, M. Ghavamzadeh, and S. Mannor, “Efficient risk-averse reinforcement learning,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 32 639–32 652, 2022.
 - [53] W. R. Clements, B. Van Delft, B.-M. Robaglia, R. B. Slaoui, and S. Toth, “Estimating risk and uncertainty in deep reinforcement learning,” *arXiv preprint arXiv:1905.09638*, 2019.
 - [54] G. B. Margolis and P. Agrawal, “Walk these ways: Tuning robot control for generalization with multiplicity of behavior,” in *Conference on Robot Learning*. PMLR, 2023, pp. 22–31.
 - [55] J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel, “High-dimensional continuous control using generalized advantage estimation,” *arXiv preprint arXiv:1506.02438*, 2015.
 - [56] S. S. Wang, “A class of distortion operators for pricing financial and insurance risks,” *Journal of risk and insurance*, pp. 15–36, 2000.
 - [57] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa *et al.*, “Isaac gym: High performance gpu-based physics simulation for robot learning,” *arXiv preprint arXiv:2108.10470*, 2021.