

Learning Risk-Aware Quadrupedal Locomotion using Distributional Reinforcement Learning

Lukas Schneider¹, Jonas Frey^{1,2}, Takahiro Miki¹, and Marco Hutter¹

Abstract—Deployment in hazardous environments requires robots to understand the risks associated with their actions and movements to prevent accidents. Despite its importance, these risks are not explicitly modeled by currently deployed locomotion controllers for legged robots. In this work, we propose a risk sensitive locomotion training method employing distributional reinforcement learning to consider safety explicitly. Instead of relying on a value expectation, we estimate the complete value distribution to account for uncertainty in the robot’s interaction with the environment. The value distribution is consumed by a risk metric to extract risk sensitive value estimates. These are integrated into Proximal Policy Optimization (PPO) to derive our method, Distributional Proximal Policy Optimization (DPPO). The risk preference, ranging from risk-averse to risk-seeking, can be controlled by a single parameter, which enables to adjust the robot’s behavior dynamically. Importantly, our approach removes the need for additional reward function tuning to achieve risk sensitivity. We show emergent risk sensitive locomotion behavior in simulation and on the quadrupedal robot ANYmal. Videos of the experiments and code are available at <https://sites.google.com/leggedrobotics.com/risk-aware-locomotion>.

I. INTRODUCTION

Legged robots can traverse rugged terrain inaccessible to wheeled or tracked systems. They can overcome sloped and uneven terrain, stones, stairs, and even large gaps by carefully choosing their foot placement. This versatility makes them ideal for operation in hazardous environments, such as cave systems [1], [2], forests [3], [4], or the surfaces of other planets [5], [6]. In these situations, the safe operation of the robot is paramount, as failure could result in hardware damage or mission failure.

Recent developments in model-free deep Reinforcement Learning (RL) [7], [8] have enabled legged robots to traverse difficult terrain reliably [9], [10]. While these systems are robust against perception failures, a notable limitation remains: they do not account for risk explicitly. Encoding behavior, such as avoiding dangerous obstacles or reducing locomotion speed on rough ground, requires adaptation of the reward formulation. Such adaptations are undesirable as learning a reliable locomotion already requires expensive and time-consuming reward function tuning.

In this work, we overcome this limitation by adapting the RL algorithm to allow risk sensitive behavior to arise without

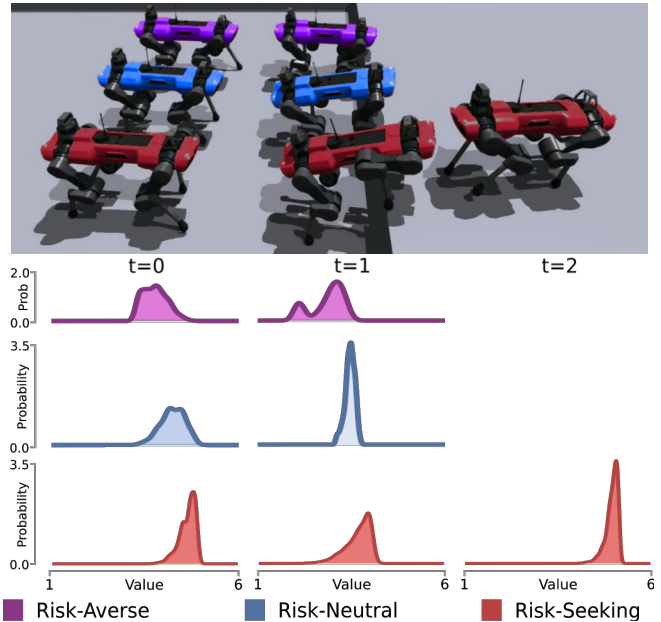


Fig. 1: Our robot learns to adapt its locomotion behavior in risky situations. When commanded to walk up a large step, the risk-averse policy (■) refuses while the risk-seeking policy (■) complies. The risk parameter, controlling the value distribution distortion, can be adapted online during deployment. Bottom: Respective risk-metric distorted value distribution per robot.

changing the reward formulation. We adopt the distributional perspective on RL [11], learning the entire distribution of returns instead of just its scalar expectation. This value distribution models the intrinsic uncertainty of the agent’s interaction with its environment. While estimating the full distribution was initially proposed for more accurate value estimates, it also captures useful information, such as the probability of catastrophic and high-return events. We use the value distribution in the actor-critic framework [12] by incorporating risk sensitive value estimates into Generalized Advantage Estimation (GAE) [13]. To extract these scalar value estimates, we employ a risk metric [14] that allows us to emphasize a specific part of the value distribution, e.g., unlikely but catastrophic events. The learned locomotion policy accounts for risks explicitly. Since the metric encodes a fixed risk preference based on a metric parameter, and no single risk preference is appropriate under all circumstances, we condition the policy on this parameter [15]. Conditioning the policy allows the adjustment of the robot’s behavior to either seek or avoid risks on demand by an operator or a high-level planner. For example, it enables the safe teleoperation of the robot by an amateur in risk-averse mode while recovering the full range of capabilities in risk-neutral / seeking mode. Figure 1 visualizes the behavior of different

¹ Robotic Systems Lab, ETH Zurich, Zurich, Switzerland

² Max Planck Institute for Intelligent Systems, Tübingen, Germany.

This work was supported by the Swiss National Science Foundation (SNSF) through project 188596, the National Centre of Competence in Research Robotics (NCCR Robotics), the European Union’s Horizon 2020 research and innovation program under grant agreement No 101016970, No 101070405, and No 852044, and an ETH Zurich Research Grant. Jonas Frey is supported by the Max Planck ETH Center for Learning Systems.

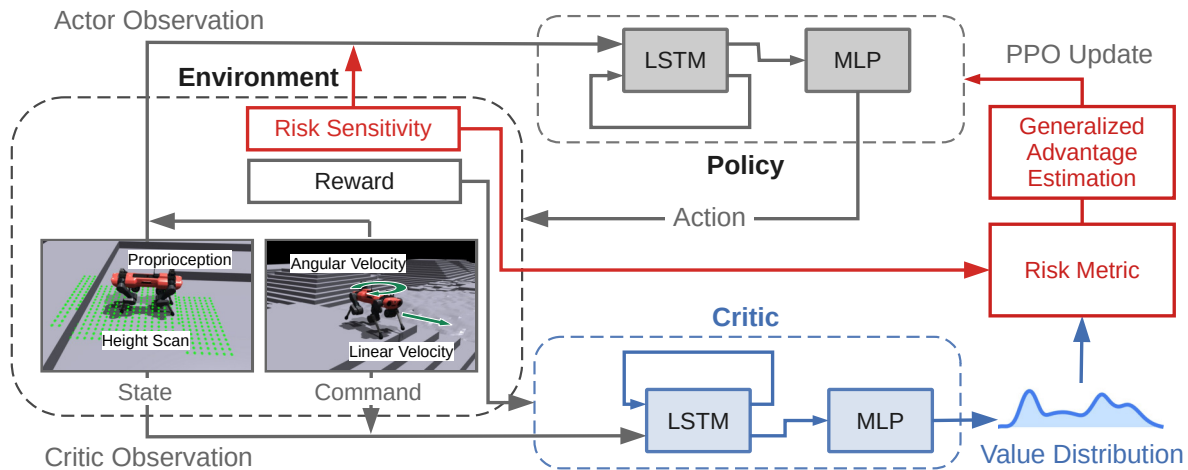


Fig. 2: Architecture overview. The critic learns to predict a value distribution, used in combination with a risk metric to update the policy. The policy is conditioned on the risk parameter. The risk parameter is part of the command, set by the operator.

risk sensitivities of a single policy.

We thoroughly analyze the risk sensitive behavior arising from our method and study the effects of using different risk metrics and parameters in simulation. We further ablate the algorithm and deploy the trained policy to the quadrupedal ANYmal robot [16], demonstrating risk sensitive behavior based on the risk parameter when climbing steps of varying heights.

Our contribution is threefold. Firstly, we integrate explicit risk modeling into the locomotion controller, a capability that currently deployed controllers do not possess. Secondly, we study the emergent risk sensitive locomotion behavior in simulation and on hardware. Lastly, on the algorithmic side, we explore introducing risk sensitivity into the RL formulation through risk sensitive advantage estimates extracted from the value distribution.

II. RELATED WORK

Legged Locomotion Legged locomotion is a thoroughly studied research area in robotics. Early successes have been achieved by model-based methods [16], [17]. However, due to their limited ability to generalize to more complex and unpredictable environments, learned methods have recently gained interest. Several deep RL algorithms have solved simulated legged locomotion tasks [7], [8], [18], [19]. Beyond that, RL policies learned in simulation have been successfully transferred to real-world scenarios, either through high simulation fidelity and data augmentation [20], [21], [10], [22], [23], [9], [24] or test-time adaptation [25].

Distributional Reinforcement Learning (Dist. RL) Dist. RL algorithms learn a value distribution, as defined by the recursive distributional Bellman Equation [11], which replaces the role of the value function in classical RL. Value-based Dist. RL algorithms [11], [26], [27], [28] use this value distribution for value-iteration, commonly maximizing its expectation when comparing different actions. These methods differ primarily in how they represent the value distribution: as a categorical distribution [11], by its cumulative distribution at fixed quantiles [26], or at arbitrary

quantiles [27], [28]. Generally, algorithms with more accurate representations achieve better results. To extend Dist. RL to continuous actions, the value distribution is used as the critic in an actor-critic architecture [12], [29], [30], [31], [32]. Several works explore Dist. RL theoretically [33], [34], [35], [36], use the additional information contained in the value distribution to improve exploration [37], or for risk sensitivity [27], [30], [15], [38], [39], [40]. Dist. RL has been used for algorithm discovery [41] and as an AI racing agent in the Gran Turismo 7 game [42]. Further, Dist. RL control policies have been successfully employed in the real world: for robotic grasping in a lab environment [43], robot soccer [44], and stratospheric balloon navigation [45].

Safe Reinforcement Learning Safe RL is a thoroughly studied research area [46]. The work of [15] is most closely related to ours, learning a risk-averse policy with online adaptation of risk-aversion. They model the value function as a Gaussian distribution and integrate it into an actor-critic architecture. When updating the policy, they maximize the Conditional Value at Risk (CVaR) metric and condition the policy on the confidence. [47] extend this work to adjust risk sensitivity automatically. Further, DSAC is a state-of-the-art risk sensitive RL algorithm that integrates a value distribution into the maximum entropy objective of SAC [48]. Several works explore risk metrics in Dist. RL with fixed risk parameters [27], [38], [39], [40], [43], [49].

III. METHOD

Our goal is to learn a locomotion policy $\pi_\phi(a|s;\beta)$ that maps the robot's state s to desired joint positions a , is conditioned on a risk parameter β , and has learned weights ϕ . The state consists of proprioceptive measurements (base velocity, projected gravity, joint positions and velocities), the previous actions, samples from a local height map around the robot, and the input commands (desired linear and angular base velocities). The risk parameter is a scalar value and depends on the metric used during training. We use the robot dynamics model from [20], including learned actuator dynamics. An overview of the architecture can be found in Figure 2.

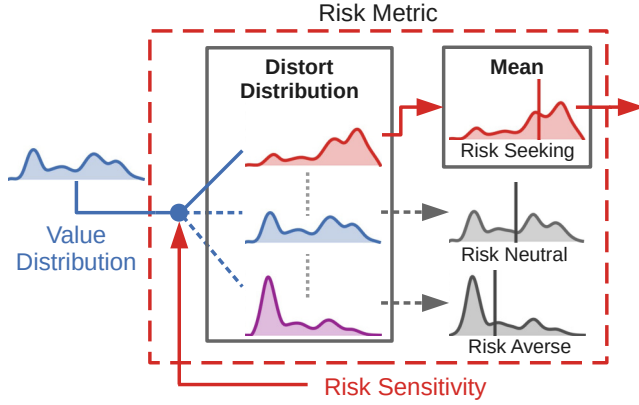


Fig. 3: Application of a risk metric. The risk sensitivity selects how the value distribution is distorted. The mean value of the distorted distribution is provided for Generalized Advantage Estimation.

Distributional Proximal Policy Optimization (DPPO)

We base our RL algorithm on the PPO algorithm [8] as it shows superior performance for learning legged locomotion in our setup, compared to DDPG [18], TD3 [50], SAC [48], and DSAC [30] (Experiment A: Section IV-B).

Critic Representation As the distributional critic, we use the Quantile Regression DQN (QR-DQN) [26]. It parameterizes the value distribution as a uniform distribution supported on $\{\theta_1(s), \dots, \theta_N(s)\}$, which can be written as

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

where δ_z is a Dirac at $z \in \mathbb{R}$. The support positions $\theta_i(s)$ are state-dependent and predicted by the critic's neural network.

Critic Updates To compute reward-to-go value targets, we rely on the sample-replacement strategy $\text{SR}(\lambda)$ [31]. This technique extends the temporal-difference $\text{TD}(\lambda)$ method [51] for computing multi-step value targets to Dist. RL. As the name suggests, it computes value targets $\mathcal{T} Z_\theta(s)$ as a mixture of sample distributions. It starts with samples of the distribution at the final time-step $Z_\theta(s_T)$. Then, going backwards in time ($t = T - 1, \dots, 0$), the sample distribution is discounted, shifted by the obtained reward r_t , and a fraction $1 - \lambda$ of samples is replaced with samples of the distribution $Z_\theta(s_t)$. 1-step and N -step target distributions are recovered by setting $\lambda = 0$ and $\lambda = 1$, respectively. Critic updates minimize the energy distance

$$\mathcal{L} = 2 \mathbb{E}_{i,j} [\theta_i - \mathcal{T} \theta_j] - \mathbb{E}_{i,j} [\mathcal{T} \theta_i - \mathcal{T} \theta_j] - \mathbb{E}_{i,j} [\theta_i - \theta_j] \quad (2)$$

between the $\text{SR}(\lambda)$ target distribution and samples from the predicted critic distribution $Z_\theta(s)$.

Rewards We adapt the rewards from [20], which primarily reward command tracking and penalize colliding / expending energy. We introduce an alive reward and modify the collision penalties to be impact force-dependent. Further, the joint motion and torque penalties are increased based on a curriculum.

Risk Quantification To quantify the risk associated with different states, we apply a risk metric to the value

distribution $Z_\theta(s)$, as depicted in Figure 3. We thereby follow [14] who argue for using distortion risk metrics, a class of metrics with properties sensible for robotics applications. Specifically, we use the Wang [52] metric, which encodes a subjective measure of risk. We further use Conditional Value at Risk (CVaR) [53], which measures the expected worst-case return (beyond a confidence level). For the distribution produced by the QR-DQN, we compute these metrics as the expectation under a distortion $g(\tau)$ of the distribution supports [30]. We denote the quantile fractions of the distribution as $(\tau_i)_{i=0, \dots, N}$ with $\tau_i = i/N$. The distorted expectation is computed as

$$V_\beta(s) = \int_0^1 g'_\beta(\tau) Z_\theta^\tau(s) d\tau = \sum_{k=1}^N (g_\beta(\tau_k) - g_\beta(\tau_{k-1})) \theta_k(s), \quad (3)$$

where we denote the τ -quantile of the value distribution as $Z_\theta^\tau(s)$. We distort the quantile fractions for CVaR as

$$g_\beta^{\text{CVaR}}(\tau) = \min\left(\frac{\tau}{\beta}, 1.0\right), \quad (4)$$

where β is a scalar risk parameter. We choose $\beta = 1$ for risk-neutrality and $0 < \beta < 1$ for risk aversion. The CVaR metric, in effect, computes the expectation of the tail of the value distribution, where β decides the cutoff. For the Wang metric, we compute the distorted quantile fractions as

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

where Φ is the standard normal distribution and β is the scalar risk parameter. We choose $\beta = 0$ for risk-neutrality, $\beta > 0$ for risk aversion, and $\beta < 0$ for risk affinity.

For the CVaR metric, we train the policy on risk parameters in the range $\beta \in (0, 1]$ while for the Wang metric, we use $\beta \in [-1.5, 1.5]$. We find that the risk sensitive behavior generalizes well beyond the parameter bounds used during training. That is, using β outside the training range results in stable but more risk-averse or affine behavior.

Policy Updates We update the policy by maximizing the usual PPO clip-objective

$$\mathcal{L} = \min \left(\frac{\pi_\phi(s|a; \beta)}{\pi_{\phi_{\text{old}}}(s|a; \beta)} A^{\pi_{\phi_{\text{old}}}(s, a; \beta)}, g(\epsilon, A^{\pi_{\phi_{\text{old}}}(s, a; \beta)}) \right) \quad (6)$$

where $g(\epsilon, A) = \begin{cases} (1 + \epsilon)A, & \text{if } A \geq 0; \\ (1 - \epsilon)A, & \text{if } A < 0. \end{cases}$

When computing the advantages $A^\pi(s_t, a_t; \beta)$, we replace the value estimates with our risk sensitive value estimates. To estimate the advantage, we use a truncated version of GAE [13] motivated by [19]. It approximates the advantage for each sample (s_t, a_t, r_t, s_{t+1}) by rolling out the policy $\pi(a|s; \beta)$ for T timesteps without looking beyond the T -th timestep. We compute truncated GAE estimates as

$$A^\pi(s_t, a_t; \beta) = \sum_{l=0}^{T-t-1} (\lambda \gamma)^l \delta_{t+l}^r \quad (7)$$

$$\text{where } \delta_t^r = r_t + \gamma V_\beta(s_{t+1}) - V_\beta(s_t)$$

and λ is the GAE bias / variance trade-off hyperparameter.

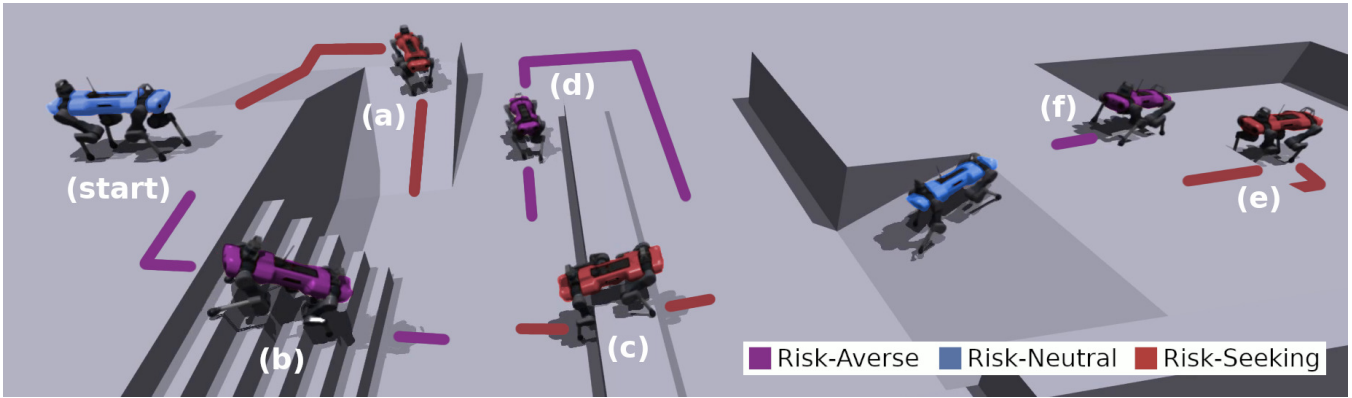


Fig. 4: Robot operated remotely along an obstacle course. Depending on the chosen path, different risk sensitivities are preferable. For the incline along path (a), a risk-seeking policy (■) can be chosen for increased walking speed. Using a risk-averse policy (■) when descending the stairs along route (b) ensures the robot’s safety. A risk-averse policy (■) won’t climb the dangerous obstacle (c) and thus would have to walk around it, along route (d). Meanwhile, setting the policy to risk-seeking (■) allows the robot to surmount the obstacle (c). To step down into the deep pit along route (e), one must set the sensitivity to risk-seeking (■). The risk-averse policy (■) will refuse to step into the pit (f) as it may lead to a crash. Video: <https://youtu.be/GGFxpF4qeVY>.

IV. EXPERIMENTS

We evaluate our method through a series of experiments. In sections IV-A to IV-C, we conduct simulation experiments to study risk sensitivity and compare our method with other approaches. In Section IV-D, we conduct hardware experiments on the ANYmal robot [16].

Training Setup To train policies, we used the legged_gym library [20], based on the IsaacGym simulator. We used the same observation and action space, except for the additional risk sensitivity. The risk sensitivity is uniformly sampled with the command during training. The training environment comprises seven different terrains with ten difficulty levels per terrain. Robots spawn at the center of each terrain with a random offset.

Simulation Experiment Setup The quantitative simulation experiments (Sections IV-B and IV-C) were conducted in a deterministic environment to enable a fair comparison. In both experiments, we spawned 72 robots at the center of each environment tile with initial z-rotations uniformly distributed across 360° and set the command to a forward velocity of 1 m/s.

A. Simulation Experiment: Obstacle Course

To demonstrate different risk sensitivities in a teleoperation scenario, we navigated a robot through the obstacle course depicted in Figure 4. The shortest path from start to goal in this environment is to descend steep stairs (b) [25cm steps], cross the barrier (c) [55cm height], walk up the ramp, and descend into the pit (e) [85cm depth]. To follow this route successfully, the operator could make use of the different risk sensitivities.

Discussion Based on risk sensitivity, the robot exhibited different locomotion behaviors. When using a risk-seeking policy, the robot generally walked faster and attempted to overcome obstacles despite the risk of falling. When commanding the robot over the barrier along the path (c), the risk-seeking policy always attempted to follow the command. In most cases, it managed to surmount the barrier but, in a few cases, got stuck on the barrier or fell down the opposite side. When sending the risk-seeking policy along path (b),

it sometimes managed to descend the stairs but fell in most cases. We obtained similar results when descending into the pit along the path (e), where the robot complied each time but often crashed. The risk-averse policy, on the other hand, demonstrated a slow but safe walking gate. It thus safely descended the stairs along path (b) on each attempt. However, the risk-averse policy refused to climb the barrier along the path (c) and descend into the pit (f). This experiment showed that different risk sensitivities change the behavior of the robot. This allows either an operator or a navigation system to make use of the risk-sensitivity to increase safety during the mission.

B. Simulation Experiment: Algorithm Performance




Fig. 5: Average return in the evaluation environment. Shaded regions indicate 95% confidence intervals across seeds and evaluation spawns. Hyperparameters for DPPO were not tuned.

We compare the attained return of our method with PPO [8], commonly used in state-of-the-art legged locomotion approaches [21], [10], and DSAC [30], a recent distributional actor-critic algorithm that can learn risk sensitive policies. DSAC has previously been tested on legged locomotion [47] and, with some adaptations, been employed in complex control scenarios [42]. For each method, we trained 10 policies. In this experiment, our method was trained with “risk-neutral” expectation instead of a risk metric. Due to computation constraints, we employed MLPs for both actor and critic across all three methods. Figure 5 shows the learning curve for each method.




Fig. 6: Average undiscounted return, fraction of episodes in which the robot terminates before episode end, and L_2 distance between commanded and actual linear velocity. PPO baseline is included for reference and does not change with risk sensitivity. Shaded regions show 95% confidence intervals across evaluation spawns. Only a single seed was trained for each method due to computation constraints. Methods are listed in table I.

Discussion We found that both DPPO (return of 27.62 after 20k iterations) and PPO (27.86) significantly outperformed DSAC (21.1). Qualitatively, DSAC, which also uses a Dist. RL formulation, failed to learn a proper locomotion policy. Despite reusing the hyperparameters from [20], extensively tuned for PPO, our method performs on-par with PPO. We thus establish our method as competitive in terms of attained return. We highlight that the core contribution of our work is not outperforming PPO in terms of attained return, but to allow for risk sensitive behavior.

C. Simulation Experiment: Risk Sensitive Performance

Name	Loss Function	Value Target
DPPO-Wang	Energy	N -step distribution [SR(λ), $\lambda = 1$]
DPPO-Wang-QH	Quantile Huber	N -step distribution
DPPO-Wang-SRL	Energy	SR(λ), $\lambda = 0.95$
DPPO-Wang-1Step	Energy	1-step distribution [SR(λ), $\lambda = 0$]

TABLE I: Overview of ablated components on DPPO-Wang.

We further compare the attained return between different ablations of our method (listed in Table I) as well as non-distributional reward-shaping baselines, introduced below. Return is computed using the original reward formulation. We show the undiscounted return in Figure 6. We also include the fraction of early terminations and the linear velocity tracking error. These two metrics underline our qualitative assessment on changing behavior with risk sensitivity from Section IV-A with quantitative results.

PPO Baseline To contrast our distributional approach, we implement PPO baselines that modulate “risk”-preference through the reward formulation. In these baselines, risk-preference is introduced before accounting for environment dynamics, setting them apart from our Dist. RL approach. Our first approach, which we call PPO1, adapts individual terms of the reward formulation. Specifically, we scale the linear velocity tracking reward by $\beta \in [0, 2]$. This scaling aims at trading-off command tracking and other rewards, most notably the alive reward. The policy should thus learn, depending on the risk parameter, to mediate differently between command compliance and keeping the robot safe.

However, reward tuning is a complex process. It is often unclear how a reward term affects the learned policy. Further, individual terms cannot be tuned in isolation as they influence one another: e.g. high torque penalties may prevent the algorithm from discovering the linear velocity tracking reward. The approach of the PPO1 baseline further increases this design complexity (especially when scaling multiple reward terms). Thus, in a second approach we call PPO2, we rescale the full reward formulation in a principled manner. We treat individual reward terms as Dirac impulses defining a probability distribution (mirroring Equation 1) and compute the distorted mean by applying the Wang metric. As risk parameter, we choose $\beta \in [-0.25, 0.25]$ as we are unable to learn a policy for a larger range.

The policy $\pi(a|s; \beta)$ learned by either baseline is conditioned on the scalar parameter β , emulating the risk parameter from our method. In both methods, the critic predicts the expected return used for GAE computation as in standard PPO. Aside from the critic, we use the same training setup, rewards, and hyperparameters as for DPPO. We expect the PPO1 baseline to learn risk-averse behavior for $\beta \in [0, 1)$ and risk-seeking behavior for $\beta \in (1, 2]$. It reverts to the initial, “risk-neutral”, reward formulation for $\beta = 1$. We expect the PPO2 baseline to exhibit behavior appropriate for the Wang metric (see Section III).

Discussion For a policy to qualify as risk sensitive, we expect the Early Terminations and the Linear Velocity Error to change appropriately across different risk settings. All methods showed such risk sensitive behavior that matches our intuition: risk-seeking policies had improved command tracking but also more premature terminations; risk-averse policies sacrificed command tracking to keep the robot safe.

The return generally increased with risk-averseness. This trend could be due to high-reward actions which, despite often failing, would still be pursued by a risk-seeking policy. In our environment, however, such risk-seeking behavior was not rewarded. Our method significantly outperformed the reward-shaping baselines regarding return. Especially the PPO2 baseline had severely lower returns. These mirrored its high early termination rate and linear velocity tracking error. For the PPO1 baseline, we experienced outlier behavior for

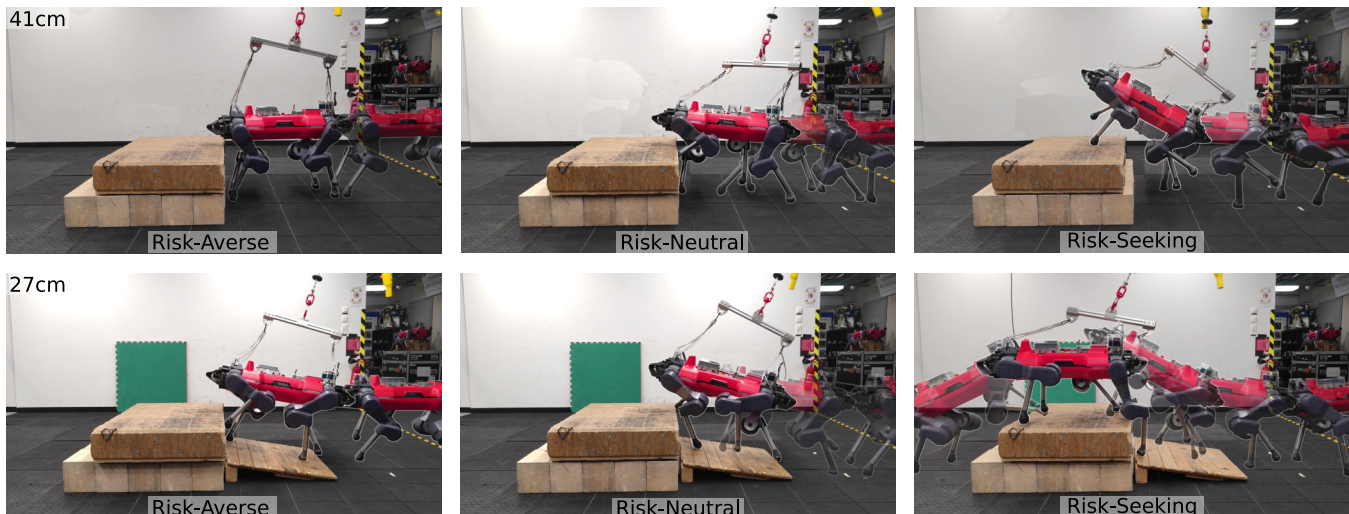


Fig. 7: ANYmal showed different step-up behaviors depending on risk sensitivity and step height. Video: <https://youtu.be/GGFxpF4qeVY>.

the most risk-averse setting: since the velocity tracking reward was set to zero during training, the robot learned to stop walking. The undesired behavior of PPO1 and low return of PPO2 demonstrate the difficulties when incorporating risk sensitivity through the standard RL framework.

Comparing our method to its ablations, none attained similar return. DPPO-CVaR most significantly underperformed in our experiments. This underperformance may be due to the CVaR metric disregarding an entire portion of the value distribution, appearing to result in less stable training. Thus, our design decisions are experimentally justified.

D. Hardware Experiment

Our learned policy transfers to the real world using appropriate domain randomization [20]. We conducted experiments on the ANYmal robot [16] with onboard perception [54] to replicate the risk sensitive behavior observed in simulation. We commanded the robot to walk up steps of height 27cm and 41cm with a speed of 1.2m/s and counted *refusals*, *failures*, and *successes* to surmount the obstacle. A high step has some risk of the robot falling and, thus, a risk-averse policy might avoid it while a risk-seeking policy should neglect the risk. The experiments are depicted in Figure 7. We collected a total of 45 tries, 8 for each policy on the 41cm step height and 7 on the 27cm step height. The results are listed in Table II.

	27cm Step			41cm Step		
Risk Sensitivity	Refusal	Failure	Succ.	Refusal	Failure	Succ.
Averse ($\beta = -1.5$)	25%	50%	25%	85.7%	14.3%	0%
Neutral ($\beta = 0.0$)	12.5%	75%	12.5%	0%	100%	0%
Seeking ($\beta = +1.5$)	0%	37.5%	62.5%	0%	100%	0%

TABLE II: Hardware experiment results. The robot received a 1.2m/s command to walk up the step. We counted the number of *refusals*, *failures* (active attempts to surmount the obstacle without ultimate success), and *successes*. The given percentages indicate how often each behavior was observed during the experiments.

Discussion The robot’s actions aligned well with the risk sensitivity of its assigned policies. When commanded to be risk-averse, the robot hesitated before climbing the step while

it tried to follow with a risk-seeking command. On the higher step height (41cm), the risk-averse policy mostly refused to walk the step. We often observed the policy standing before the obstacle, unresponsive to its forward command. The risk-seeking and neutral policies, while failing to surmount the obstacle as well, exhibited more aggressive behavior. The difference between the risk-neutral and risk-seeking policies became emphasized for the smaller step size (27cm). In this experiment, the risk-seeking policy complied every time, often surmounting the obstacle. The risk-neutral policy mostly attempted to overcome the obstacle but failed. A slight inconsistency arose between the risk-neutral and risk-averse policies, where the latter succeeded in walking the step more often. Overall, the risk-neutral policy attempted to walk the step more often than the risk-averse policy. These observations mirror our findings in simulation and confirm that risk sensitivity successfully transfers to hardware.

V. CONCLUSION

Our work presents a Dist. RL approach for learning risk sensitive locomotion policies. These policies allow the robot to adapt its behavior to environment risks. A risk preference is encoded through a single parameter that can be modulated during deployment. To incorporate risk sensitivity into the RL formulation, we introduce risk sensitive advantage estimates. The emergent risk sensitive behavior, demonstrated in simulation and on hardware, opens up a new application domain for Dist. RL in real-world applications. Our work directly benefits teleoperation and navigation, given that a user / planner can now change the risk affinity of the robot during deployment, increasing safety while retaining full locomotion capability. Further research is needed to establish procedures on evaluating risk sensitivity and study risk sensitive Dist. RL from a theoretical standpoint: how the choice of risk metric impacts the resulting policy and how well the estimated value distribution matches the true return distribution are some suggestions. Additionally, integrating DPPO into a navigation system that utilizes risk sensitivity is promising for future work.

REFERENCES

- [1] A. Bouman, M. F. Ginting, N. Alatur, M. Palieri, D. D. Fan, T. Touma, T. Pailevanian, S.-K. Kim, K. Otsu, J. W. Burdick, and A. Akbar Agha-mohammadi, "Autonomous Spot: Long-Range Autonomous Exploration of Extreme Environments with Legged Locomotion," *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2518–2525, 2020.
- [2] M. Tranzatto, M. Dharmadhikari, L. Bernreiter, M. Camurri, S. Khatkhat, F. Mascari, P. Pfreundschuh, D. Wisth, S. Zimmermann, M. Kulkarni, V. Reijgwart, B. Casseau, T. Homberger, P. D. Petris, L. Ott, W. Tubby, G. Waibel, H. Nguyen, C. Cadena, R. Buchanan, L. Wellhausen, N. Khedekar, O. Andersson, L. Zhang, T. Miki, T. Dang, M. Mattamala, M. Montenegro, K. Meyer, X. Wu, A. Briod, M. Mueller, M. Fallon, R. Siegwart, M. Hutter, and K. Alexis, "Team CERBERUS Wins the DARPA Subterranean Challenge: Technical Overview and Lessons Learned," 2022.
- [3] J. Frey, D. Hoeller, S. Khattak, and M. Hutter, "Locomotion Policy Guided Traversability Learning using Volumetric Representations of Complex Environments," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022.
- [4] J. Frey, M. Mattamala, N. Chebroul, C. Cadena, M. Fallon, and M. Hutter, "Fast Traversability Estimation for Wild Visual Navigation," in *Proceedings of Robotics: Science and Systems*, 2023.
- [5] A. Dettmann, S. Plathaber, V. Bargsten, R. Dominguez, G. Cerilli, M. Marchitto, G. Fink, M. Focchi, V. Barasuol, C. Semini, and R. Marc, "Towards a generic navigation and locomotion control system for legged space exploration." 16th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA 2022), 2022.
- [6] H. Kolvenbach, M. Breitenstein, C. Gehring, and M. Hutter, "Scalability Analysis of Legged Robots for Space Exploration." 68th International Astronautical Congress (IAC), 2017.
- [7] 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.
- [8] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal Policy Optimization Algorithms." *CoRR*, vol. abs/1707.06347, 2017.
- [9] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning quadrupedal locomotion over challenging terrain," *Science Robotics*, vol. 5, no. 47, oct 2020.
- [10] T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning robust perceptive locomotion for quadrupedal robots in the wild." *Science Robotics*, vol. 7, no. 62, p. eabk2822, 2022.
- [11] M. G. Bellemare, W. Dabney, and R. Munos, "A Distributional Perspective on Reinforcement Learning," in *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, ser. Proceedings of Machine Learning Research, vol. 70. PMLR, 2017, pp. 449–458.
- [12] G. Barth-Maron, M. W. Hoffman, D. Budden, W. Dabney, D. Horgan, D. TB, A. Muldal, N. Heess, and T. P. Lillicrap, "Distributed Distributional Deterministic Policy Gradients," in *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.
- [13] J. Schulman, P. Moritz, S. Levine, M. I. Jordan, and P. Abbeel, "High-Dimensional Continuous Control Using Generalized Advantage Estimation," in *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.
- [14] A. Majumdar and M. Pavone, "How Should a Robot Assess Risk? Towards an Axiomatic Theory of Risk in Robotics," in *Robotics Research, The 18th International Symposium, ISRR 2017, Puerto Varas, Chile, December 11-14, 2017*, ser. Springer Proceedings in Advanced Robotics, vol. 10. Springer, 2017, pp. 75–84.
- [15] Y. C. Tang, J. Zhang, and R. Salakhutdinov, "Worst Cases Policy Gradients," in *3rd Annual Conference on Robot Learning, CoRL 2019, Osaka, Japan, October 30 - November 1, 2019, Proceedings*, ser. Proceedings of Machine Learning Research, vol. 100. PMLR, 2019, pp. 1078–1093.
- [16] M. Hutter, C. Gehring, D. Jud, A. Lauber, C. D. Bellicoso, V. Tsounis, J. Hwangbo, K. Bodie, P. Fankhauser, M. Bloesch, R. Diethelm, S. Bachmann, A. Melzer, and M. Hoepflinger, "ANYmal - a highly mobile and dynamic quadrupedal robot," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016, pp. 38–44.
- [17] G. Bledt, M. J. Powell, B. Katz, J. Di Carlo, P. M. Wensing, and S. Kim, "MIT Cheetah 3: Design and Control of a Robust, Dynamic Quadruped Robot," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 2245–2252.
- [18] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," in *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.
- [19] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. P. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu, "Asynchronous Methods for Deep Reinforcement Learning," in *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, ser. JMLR Workshop and Conference Proceedings, vol. 48. JMLR.org, 2016, pp. 1928–1937.
- [20] N. Rudin, D. Hoeller, P. Reist, and M. Hutter, "Learning to Walk in Minutes Using Massively Parallel Deep Reinforcement Learning," in *5th Annual Conference on Robot Learning*, 2021.
- [21] A. Agarwal, A. Kumar, J. Malik, and D. Pathak, "Legged Locomotion in Challenging Terrains using Egocentric Vision," in *Conference on Robot Learning, CoRL 2022, 14-18 December 2022, Auckland, New Zealand*, ser. Proceedings of Machine Learning Research, vol. 205. PMLR, 2022, pp. 403–415.
- [22] Z. Xie, X. Da, M. van de Panne, B. Babich, and A. Garg, "Dynamics Randomization Revisited: A Case Study for Quadrupedal Locomotion," *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4955–4961, 2020.
- [23] J. Tan, T. Zhang, E. Coumans, A. Iscen, Y. Bai, D. Hafner, S. Bohez, and V. Vanhoucke, "Sim-to-Real: Learning Agile Locomotion For Quadruped Robots," in *Robotics: Science and Systems XIV, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA, June 26-30, 2018*, 2018.
- [24] 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. eaau5872, 2019.
- [25] L. Smith, J. C. Kew, X. B. Peng, S. Ha, J. Tan, and S. Levine, "Legged Robots that Keep on Learning: Fine-Tuning Locomotion Policies in the Real World," *2022 International Conference on Robotics and Automation (ICRA)*, pp. 1593–1599, 2021.
- [26] W. Dabney, M. Rowland, M. G. Bellemare, and R. Munos, "Distributional Reinforcement Learning With Quantile Regression," in *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18), the 30th Innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*. AAAI Press, 2018, pp. 2892–2901.
- [27] W. Dabney, G. Ostrovski, D. Silver, and R. Munos, "Implicit Quantile Networks for Distributional Reinforcement Learning," in *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, ser. Proceedings of Machine Learning Research, vol. 80. PMLR, 2018, pp. 1104–1113.
- [28] D. Yang, L. Zhao, Z. Lin, T. Qin, J. Bian, and T. Liu, "Fully Parameterized Quantile Function for Distributional Reinforcement Learning," in *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, 2019, pp. 6190–6199.
- [29] A. Kuznetsov, P. Shvechikov, A. Grishin, and D. P. Vetrov, "Controlling Overestimation Bias with Truncated Mixture of Continuous Distributional Quantile Critics," in *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, ser. Proceedings of Machine Learning Research, vol. 119. PMLR, 2020, pp. 5556–5566.
- [30] X. Ma, Q. Zhang, L. Xia, Z. Zhou, J. Yang, and Q. Zhao, "Distributional Soft Actor Critic for Risk Sensitive Learning," *CoRR*, vol. abs/2004.14547, 2020.
- [31] D. W. Nam, Y. Kim, and C. Y. Park, "GMAC: A Distributional Perspective on Actor-Critic Framework," in *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, ser. Proceedings of Machine Learning Research, vol. 139. PMLR, 2021, pp. 7927–7936.
- [32] R. Singh, K. Lee, and Y. Chen, "Sample-based Distributional Policy Gradient," in *Learning for Dynamics and Control Conference, L4DC 2022, 23-24 June 2022, Stanford University, Stanford, CA, USA*, ser.

- Proceedings of Machine Learning Research, vol. 168. PMLR, 2022, pp. 676–688.
- [33] M. G. Bellemare, I. Danihelka, W. Dabney, S. Mohamed, B. Lakshminarayanan, S. Hoyer, and R. Munos, “The Cramer Distance as a Solution to Biased Wasserstein Gradients,” *CoRR*, vol. abs/1705.10743, 2017.
- [34] K. Sun, Y. Zhao, Y. Liu, E. Shi, Y. Wang, X. Yan, B. Jiang, and L. Kong, “Interpreting Distributional Reinforcement Learning: A Regularization Perspective,” 2021.
- [35] M. Rowland, R. Dadashi, S. Kumar, R. Munos, M. G. Bellemare, and W. Dabney, “Statistics and Samples in Distributional Reinforcement Learning,” in *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, ser. Proceedings of Machine Learning Research, vol. 97. PMLR, 2019, pp. 5528–5536.
- [36] M. Rowland, R. Munos, M. G. Azar, Y. Tang, G. Ostrovski, A. Harutyunyan, K. Tuyls, M. G. Bellemare, and W. Dabney, “An Analysis of Quantile Temporal-Difference Learning,” 2023.
- [37] B. Mavrin, S. Zhang, H. Yao, L. Kong, K. Wu, and Y. Yu, “Distributional Reinforcement Learning for Efficient Exploration.” *CoRR*, vol. abs/1905.06125, 2019.
- [38] J. Bernhard, S. Pollok, and A. Knoll, “Addressing Inherent Uncertainty: Risk-Sensitive Behavior Generation for Automated Driving using Distributional Reinforcement Learning,” in *2019 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, jun 2019.
- [39] S. Stanko and K. Macek, “Risk-averse Distributional Reinforcement Learning: A CVaR Optimization Approach,” in *International Joint Conference on Computational Intelligence*, 2019.
- [40] T. Théate and D. Ernst, “Risk-Sensitive Policy with Distributional Reinforcement Learning,” 2022.
- [41] A. Fawzi, M. Balog, A. Huang, T. Hubert, B. Romera-Paredes, M. Barekatin, A. Novikov, F. Ruiz, J. Schrittwieser, G. Swirszcz, D. Silver, D. Hassabis, and P. Kohli, “Discovering faster matrix multiplication algorithms with reinforcement learning,” *Nature*, vol. 610, pp. 47–53, 10 2022.
- [42] P. Wurman, S. Barrett, K. Kawamoto, J. MacGlashan, K. Subramanian, T. Walsh, R. Capobianco, A. Devlic, F. Eckert, F. Fuchs, L. Gilpin, P. Khandelwal, V. Kompella, H. Lin, P. MacAlpine, D. Oller, T. Seno, C. Sherstan, M. Thomure, and H. Kitano, “Outracing champion Gran Turismo drivers with deep reinforcement learning,” *Nature*, vol. 602, pp. 223–228, 02 2022.
- [43] C. Bodnar, A. Li, K. Hausman, P. Pastor, and M. Kalakrishnan, “Quantile QT-Opt for Risk-Aware Vision-Based Robotic Grasping,” in *Robotics: Science and Systems XVI*. Robotics: Science and Systems Foundation, jul 2020.
- [44] T. Haarnoja, B. Moran, G. Lever, S. H. Huang, D. Tirumala, M. Wulfmeier, J. Humplik, S. Tunyasuvunakool, N. Y. Siegel, R. Hafner, M. Bloesch, K. Hartikainen, A. Byravan, L. Hasenclever, Y. Tassa, F. Sadeghi, N. Batchelor, F. Casarini, S. Saliceti, C. Game, N. Sreendra, K. Patel, M. Gwira, A. Huber, N. Hurley, F. Nori, R. Hadsell, and N. Heess, “Learning Agile Soccer Skills for a Bipedal Robot with Deep Reinforcement Learning,” 2023.
- [45] M. Bellemare, S. Candido, P. Castro, J. Gong, M. Machado, S. Moitra, S. Ponda, and Z. Wang, “Autonomous navigation of stratospheric balloons using reinforcement learning,” *Nature*, vol. 588, pp. 77–82, 12 2020.
- [46] J. García, Fern, and o Fernández, “A Comprehensive Survey on Safe Reinforcement Learning,” *Journal of Machine Learning Research*, vol. 16, no. 42, pp. 1437–1480, 2015.
- [47] F. Schubert, T. Eimer, B. Rosenhahn, and M. Lindauer, “Automatic Risk Adaptation in Distributional Reinforcement Learning,” 2021.
- [48] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, “Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor,” in *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, ser. Proceedings of Machine Learning Research, vol. 80. PMLR, 2018, pp. 1856–1865.
- [49] N. A. Urpí, S. Curi, and A. Krause, “Risk-Averse Offline Reinforcement Learning,” 2021.
- [50] S. Fujimoto, H. van Hoof, and D. Meger, “Addressing Function Approximation Error in Actor-Critic Methods,” in *ICML*, ser. Proceedings of Machine Learning Research, vol. 80. PMLR, 2018, pp. 1582–1591.
- [51] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. The MIT Press, 2018.
- [52] S. S. Wang, “A Class of Distortion Operators for Pricing Financial and Insurance Risks,” *The Journal of Risk and Insurance*, vol. 67, no. 1, pp. 15–36, 2000.
- [53] R. T. Rockafellar and S. Uryasev, “Optimization of conditional value-at risk,” *Journal of Risk*, vol. 3, pp. 21–41, 2000.
- [54] T. Miki, L. Wellhausen, R. Grandia, F. Jenelten, T. Homberger, and M. Hutter, “Elevation Mapping for Locomotion and Navigation using GPU,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2022, Kyoto, Japan, October 23-27, 2022*. IEEE, 2022, pp. 2273–2280.