

Roll-Drop: accounting for observation noise with a single parameter

Luigi Campanaro
Daniele De Martini
Siddhant Gangapurwala
Wolfgang Merkt
Ioannis Havoutis

LUIGI@ROBOTS.OX.AC.UK
 DANIELE@ROBOTS.OX.AC.UK
 SIDDHANT@ROBOTS.OX.AC.UK
 WOLFGANG@ROBOTS.OX.AC.UK
 IOANNIS@ROBOTS.OX.AC.UK

Department of Engineering Science, University of Oxford

Editors: N. Matni, M. Morari, G. J. Pappas

Abstract

This paper proposes a simple strategy for sim-to-real in Deep-Reinforcement Learning (DRL) – called Roll-Drop – that uses dropout during simulation to account for observation noise during deployment without explicitly modelling its distribution for each state. DRL is a promising approach to control robots for highly dynamic and feedback-based manoeuvres, and accurate simulators are crucial to providing cheap and abundant data to learn the desired behaviour. Nevertheless, the simulated data are noiseless and generally show a distributional shift that challenges the deployment on real machines where sensor readings are affected by noise. The standard solution is modelling the latter and injecting it during training; while this requires a thorough system identification, Roll-Drop enhances the robustness to sensor noise by tuning only a single parameter. We demonstrate an 80% success rate when up to 25% noise is injected in the observations, with twice higher robustness than the baselines. We deploy the controller trained in simulation on a Unitree A1 platform and assess this improved robustness on the physical system. Additional resources at: <https://sites.google.com/oxfordrobotics.institute/roll-drop>

Keywords: Sim-to-real; Legged Locomotion; Reinforcement Learning.

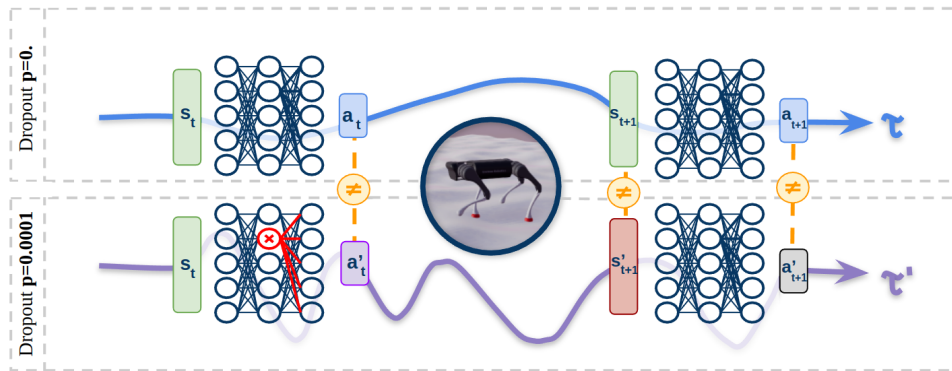


Figure 1: We present two policies, $\hat{\pi}$ trained with dropout during rollout (bottom) and π without it (top). As the time-step t precedes the occurrence of the first dropout and the training adopts the same random seed, both policies visited the same states and actions. After the first dropout is triggered, the policies will follow different trajectories, \mathcal{T}_π and $\mathcal{T}'_{\hat{\pi}}$: $[a_t, \dots, a_T]_\pi \neq [a'_t, \dots, a'_T]_{\hat{\pi}}$, $[r_t, \dots, r_T]_\pi \neq [r'_t, \dots, r'_T]_{\hat{\pi}}$, and $[s_{t+1}, \dots, s_T]_\pi \neq [s'_{t+1}, \dots, s'_T]_{\hat{\pi}}$. This alters the visited states and prevents high sensitivity of policies to noiseless observations.

1. Introduction

Deep-Reinforcement Learning (DRL) gained traction in the legged robotics community as a promising approach to the control problem, enabling highly dynamic and sophisticated locomotion capabilities (Lee et al., 2019; Yang et al., 2020; Kumar et al., 2021). The sample complexity associated with high-dimensional problems such as locomotion, the risk of breaking the machines at the beginning of the training and the difficulty of resetting the robots make the use of physics simulators (Hwangbo et al., 2018; Makoviychuk et al., 2021) appealing for training DRL control policies. However, this convenience often requires addressing the *reality gap* between the simulated training and physical deployment domains.

Strategies to address such a reality gap include accurately identifying properties such as Center of Mass (CoM), mass and inertia of robot links, impedance gains, system communication delays, friction, and actuation dynamics (Hwangbo et al., 2019; Lee et al., 2020). In addition, relevant distributions suitable for domain randomisation need to be selected (Tan et al., 2018; Lee et al., 2019); as part of such randomisation of the environment, sensory noise needs modelling and it is introduced in simulation during training (Jakobi et al., 1995; Hwangbo et al., 2019).

We recently proposed Extended Random Force Injection (ERFI) (Campanaro et al., 2022) to handle system and actuation uncertainty as an alternative to a complete system and distribution identification for dynamics randomisation. We demonstrate state-of-the-art sim-to-real performances by only randomising (and tuning) two parameters. However, the robustness showed by ERFI in challenging conditions did not explicitly encompass modelling noise in observations.

In this work, we propose Roll-Drop, a method that improves the robustness of DRL-based locomotion controllers to observation noise by introducing dropout during rollout. In continuation with ERFI’s simplicity, Roll-Drop only needs tuning a single parameter.

In the following sections, we present the method, analyse the results, and compare the robustness of alternatives to the injection of noise in the state space of the policy. Roll-Drop demonstrates an 80% success rate when up to 25% noise is injected in the observations, whereas in the same conditions other techniques experienced less than 40% success rate. The policies were trained in simulation on flat ground and deployed on a Unitree A1.

2. Related Work

Modern robots are equipped with diverse sensors to ensure acceptable levels of autonomy by estimating either the robot’s state or the surrounding environment. Such sensors include Inertial Measurement Units (IMUs), joint encoders (Hubicki et al., 2016), torque and contact sensors (Hutter et al., 2016), RGBD cameras (Rudin et al., 2022; Gangapurwala et al., 2022; Miki et al., 2022), and lidar scanners (Mattamala et al., 2022). DRL approaches applied to locomotion controllers conveniently train policies that can take advantage of such rich sensory information.

Simulators are paramount here to reducing costs and training time while ensuring safety during the delicate training procedure. Moreover, simulators provide the repeatability necessary to investigate eventual undesired behaviour. However, in contrast to real sensors, simulators provide perfect and noiseless information far from what the policy would experience when deployed on a real robot, causing an additional sim-to-real gap to be addressed.

Research has focused on tackling the mismatch between simulated and real sensors by directly modelling the noise from real systems and injecting it into the network state during training. Hwangbo et al. (2019) sample the joint velocity noise from uniform distributions, similarly

to Lee et al. (2020) for linear and angular velocity noise; these were then added to the simulator’s observations to improve robustness. Bohez et al. (2022) use instead normal distributions to model observation noise for joint positions, angular velocity, linear acceleration, and base orientation, while Siekmann et al. (2021); Yu et al. (2022) also include the joint-encoder offsets, which were sampled from a uniform distribution. Gangapurwala et al. (2022); Miki et al. (2022), instead, focus on exteroceptive sensors and inject noise into the height maps to foster the controller robustness to artefacts and sudden spikes.

Additionally, determining the noise characteristics is a delicate and costly process. Often little detail on the process is provided, and ablation studies supporting the necessity of such randomisation are absent. Roll-Drop addresses this lack of information as one parameter is enough to characterise the implementation.

3. Problem Definition: Blind Quadrupedal Locomotion

We model a quadrupedal system as a floating base B described by the reference frame \mathcal{B} , represented w.r.t. a world reference frame \mathcal{W} , whose z -axis aligns with the gravity axis. \mathcal{B} ’s x -axis $x_{\mathcal{B}}$ points in the forward direction of motion of B , the y -axis $y_{\mathcal{B}}$ to the left and the z -axis $z_{\mathcal{B}}$ upwards. The base position is then expressed as $r_B \in \mathbb{R}^3$, and the orientation, $q_B \in SO(3)$, is represented by a unit quaternion, whose corresponding rotation matrix is denoted as $\mathbf{R}_B \in SO(3)$.

In this work, we will employ a Unitree A1 quadruped, whose four legs are composed of three joints each. We will refer to the front-right leg as FR, to the front-left leg as FL, to the hind-right as HR, and to the hind-left leg as HL. Each leg has a hip adduction/abduction HAA, hip flexion/extension HFE, and knee flexion/extension KFE joint. For example, we refer to the front-right hip flexion/extension as FR_HFE. The vector $q_j \in \mathbb{R}^{n_j}$ – in our system, $n_j = 12$ – contains the angular positions of the rotational joints of all limbs, which are actuated through an impedance control, simplified as described by Peng and van de Panne (2017):

$$\Gamma_j = K_p(q_j^* - q_j) - K_d\dot{q}_j. \quad (1)$$

where Γ_j are the actuation torques on the joints, q_j^* is the vector representing desired joint positions, and K_p and K_d refer to the position and velocity tracking gains, respectively, which in our system are $K_p = 15.0$ and $K_d = 1.0$.

3.1. Reinforcement Learning

The Reinforcement Learning (RL) problem is modelled as an Markov Decision Process (MDP) including a state space \mathbf{S} , an action space \mathbf{A} , an initial state distribution $p_1(s_1)$, a transition dynamics $p(s_{t+1}|s_t, a_t)$ compliant with the Markov property $p(s_{t+1}|s_1, a_1, \dots, s_t, a_t) = p(s_{t+1}|s_t, a_t)$ for any trajectory $\mathcal{T}_{1:t} = [(s_1, a_1, r_1), (s_2, a_2, r_2), \dots, (s_t, a_t, r_t)]$, where $r_i = R(s_i, a_i)$ is the reward obtained from a reward function $R : \mathbf{S} \times \mathbf{A} \rightarrow \mathbb{R}$. In all the previous, $s_i \in \mathbf{S}$ and $a_i \in \mathbf{A}$.

A policy – in our case, the controller – selects actions in the MDP given a specific state. The policy – denoted by π_θ , where $\theta \in \mathbb{R}^n$ is a vector of n parameters – is stochastic, and $\pi_\theta(a_t|s_t)$ is the conditional probability density of a_t associated with the policy. The agent uses its policy to interact with the MDP, realising the trajectory of states, actions, and rewards $\mathcal{T}_{1:T} = (s_1, a_1, r_1), \dots, (s_T, a_T, r_T)$.

Table 1: (a) PPO hyper-parameters used for training the Unitree A1 policy; (b) the rewards adopted during training, and their weights.

(a)		(b)		
Hyperparameter	Value	Definition	Weight	
Control dt	0.02 [s]	Base orientation	$k_c \cdot \ \mathbf{R}_B^z - [0, 0, 1]\ ^2$	-30
Sim dt	0.002 [s]	Base linear velocity	$\phi(v_{b_x,y}^*, v_{b_x,y}, 5)$	15
Batch size	25600	Base angular velocity	$\phi(v_{b_z}^*, \omega_{b_z}, 5)$	15
Mini-batch size	6400	Action smoothness	$k_c \cdot \ \mathbf{q}_{j_t}^* - \mathbf{q}_{j_{t-1}}^*\ ^2$	-7
Number of epochs	8	Feet clearance	$k_c \cdot \sum_{n=0}^{n < 4} (0.1 - f_{z_n})^2$	-400
Clip range	0.2	Feet sleep	$k_c \cdot \ \dot{f}_{x,y}\ ^2$	-8
Entropy coefficient	0.	Joint position	$k_c \cdot \ \mathbf{q}_j - \mathbf{q}_j^N\ ^2$	-4
Discount factor	0.996	Joint velocity	$k_c \cdot \ \dot{\mathbf{q}}_j\ ^2$	-0.01
GAE discount factor	0.95	Joint torque	$k_c \cdot \ \boldsymbol{\tau}_j\ ^2$	-0.4
Learning rate	$1e^{-4}$	Feet swing duration	$\sum_{n=0}^3 (\mathbf{t}_{air,n} - 0.5)$	8
		Pronking gait	$k_c \cdot \sum_{n=0}^3 (f_{c_n} \cdot 1)$	-35

The policy π is trained through an optimisation problem to maximise the cumulative discounted reward it obtained from the starting state, expressed as $\pi^* = \operatorname{argmax} \mathbb{E}[r_1^\gamma | \pi]$, where r_t^γ is the total discounted reward from time-step t onward, as $r_t^\gamma = \sum_{k=t}^T \gamma^{k-t} r(s_k, a_k)$, where $0 < \gamma < 1$.

3.2. Implementation

The quadruped robot is required to follow a velocity command $s_c = [v_x, v_y, \psi]_{\mathcal{B}}$ on flat ground using proprioceptive information. Here v_x and v_y are the linear velocities along $x_{\mathcal{B}}$ and $y_{\mathcal{B}}$ respectively, while ψ is the angular velocity around $z_{\mathcal{B}}$.

The state is represented as $s := \langle s_r, s_v, s_{j_p}, s_{j_v}, s_a, s_f, s_c \rangle \in \mathbb{R}^{196}$, where $s_r^{\mathcal{B}} \in \mathbb{R}^3$ is the last row of the rotation matrix \mathbf{R}_B , $s_v \in \mathbb{R}^6$ is the base linear and angular velocities, $s_{j_p}^{\mathcal{B}} \in \mathbb{R}^{84}$ is the history of joint position errors and $s_{j_v}^{\mathcal{B}} \in \mathbb{R}^{84}$ is the history of joint velocities, $s_a \in \mathbb{R}^{12}$ is the previous action, $s_f \in \mathbb{R}^4$ is the contact state of the feet, and $s_c^{\mathcal{B}} \in \mathbb{R}^3$ is the velocity command. The actions $a \in \mathbb{R}^{12}$ are retrieved from the policy π – implemented as a Multi-Layer Perceptron (MLP) formed by three layers of size [512, 256, 256] – and interpreted as the reference joint positions \mathbf{q}_j^* , tracked by the impedance controller in Equation (1). The onboard state estimator does not provide the base linear velocity in s_v ; hence, we estimate it and s_f similarly to Ji et al. (2022) through an MLP of size [128, 128, 128]. We train π on flat ground using Proximal Policy Optimization (PPO) (Schulman et al., 2017) until convergence (Figure 7(b)), adopting the rewards and hyper-parameters in Table 1(a).

4. Roll-Drop

The proposed method, Roll-Drop, exploits the concept of using dropout to mimic an observation noise to improve the network’s robustness in a sim-to-real deployment scenario. In particular, Roll-Drop adds a customised dropout layer (Hinton et al., 2012), active only during rollouts and turned off during training. The resulting random perturbations (as shown in Figure 1) cause the policy π to explore regions of the state space \mathbf{s} and action space \mathbf{A} different from the standard training, as in Figure 5.




Figure 2: (a) Evaluation environment: the non-flat terrain is more realistic and it brings stochasticity to the evaluation of the robustness to observation noise. (b) Default settings for the testing environment.

When the dropout is not present, the actions are sampled with a policy $\pi_\theta(a|s)$: $\pi_\theta(a|s) = \mu_\theta(s) + \mathcal{N}(0, \sigma)$, where $\mu_\theta(s)$ is the output of the network. When, instead, dropout is included the parameters θ become $\hat{\theta} = \theta + \delta\theta$, and consequently $\pi_{\theta+\delta\theta}(a|s) = \mu_{\theta+\delta\theta}(s) + \mathcal{N}(0, \sigma)$. Based on this, in a state s : $\mu_\theta(s) \mapsto \alpha$, while $\mu_{\theta+\delta\theta}(s) \mapsto \hat{\alpha}$ with $\hat{\alpha} = \alpha + \delta\alpha$ and $\delta\alpha$ function of $\hat{\theta}$.

Assuming deterministic dynamics and same initialisation, the transition probability can be reformulated as the transition function $\mathcal{P}(s, \alpha) \rightarrow \zeta$, when dropout is inactive, and $\mathcal{P}(s, \hat{\alpha}) \rightarrow \hat{\zeta}$ otherwise, where ζ and $\hat{\zeta}$ are the next states. Similarly to $\hat{\alpha}$, $\hat{\zeta}$ can be expressed as $\hat{\zeta} = \zeta + \delta\zeta$, where $\delta\zeta = f(\alpha + \delta\alpha)$.

At the next time-step ($t + 1$), when dropout is inactive we can expect $\mathcal{P}(s', a') \rightarrow s''$, whereas when dropout is active $\mathcal{P}(s' + \delta s', a' + \delta a') \rightarrow s'' + \delta s''$. Here $\delta s'$ and $\delta s''$ represent the discrepancy between the transitions happening adopting $\pi_\theta(a|s)$ and $\pi_{\theta+\delta\theta}(a|s)$.

In this work, we added a single layer of Roll-Drop after the second layer of the MLP network. Notably, since the dropout-injected noise happens only during rollout, π develops reflexes to recover from dangerous states and becomes more robust to perturbations; conversely, adding dropouts during training does not allow the policy to develop reactions to perturbations.

4.1. Tuning Roll-Drop probability

Similarly to other randomisation techniques (Tobin et al., 2017; Valassakis et al., 2020; Campanaro et al., 2022), the tuning of the Roll-Drop probability is carried out empirically: At first, the environment (defined in Tables 1(a) and 1(b)) is tuned for tracking a velocity command s_c on flat ground **without** any randomisation and using a fixed random seed. After the policy converges to the desired behaviour, the Roll-Drop layer is included in the network, and the dropout probability is increased (starting from $p = 0$.) until the training is stable again. This can be seen in Figure 3(a), where we tested different dropout probabilities and how they affected the training convergence.

5. Experimental Setup

To assess the performance of the method proposed we run several experiments with different levels of noise affecting the observations. The environment’s settings are fixed as in Figure 2(b), the robot is commanded a constant velocity $s_c = [v_x, v_y, v_z]$, where $v_x = 0.5[m/s]$ is the only non zero component.

Alongside these settings we included a mild rough terrain to better represent realistic conditions, as in Figure 2(a). In the environment defined as above we varied the amount of noise (n) from 0% to 60% as in Equation (2), with a step of 5%, and 100 experiments were run for each noise configuration (randomising the spawning point of the robot on the rough terrain).

$$s_t \leftarrow s_t + n \cdot \mathcal{U}(-1, 1) \cdot s_t, \text{ where } n \in [0, 0.6] \quad (2)$$

The success rate in Figure 7(a) is measured across the 100 experiments carried out for each percentage of injected noise. To successfully complete the evaluation the robot does not have to fall on the ground and it has to walk for at least 1 [m] in the direction of the velocity commanded, if one of the two conditions is not respected the experiment is considered a failure. The ratio between the successful runs and the total number runs gives the success rate.

6. Results and Discussion

We compared Roll-Drop ($p = 0.0001$) against *No Randomisation* –which is based on the original environment used for Roll-Drop but without dropout, in Section 4.1–, against ERFI (Campanaro et al., 2022) that demonstrated state-of-the-art robustness to external perturbations, against dropout during training ($p = 0.001$), and finally a mixture of dropout during training ($p = 0.001$) plus dropout during rollout ($p = 0.0001$). From the results in Figure 7(a), the most robust method to the injection of noise in the observations is Roll-Drop, which retained 80% success rate when more than 25% of noise was injected. The performance of the policies trained with other techniques degrades quickly as soon as noise is injected, suggesting strong sensitivity to observation distribution encountered during training. Note that all the controllers were trained and tested adopting the same random seed.

6.1. Dropout during training

We motivate the adoption of dropouts during rollouts (Roll-Drop) in Section 4, nonetheless we investigated the performances resulting from adopting dropouts during training, and during both training and rollouts. This is depicted in Figure 3(b), where we show the effects of different dropout probabilities when it is applied during training and not during rollout. A cluster of lines can be identified with dropout probability $\in [0.01, 0.1]$, their maximum reward oscillates around 0.3, which corresponds to the robot standing still. Based on our experience, the randomness injected by dropout does not allow the network to correlate inputs and outputs well, and by standing still the policy avoids the termination reward (when the robot falls on the ground), while still receiving some positive points from the rewards in Table 1(b). Indeed, as soon as the dropout probability is lowered to $p = 0.001$ the total reward increases, and the robot starts walking again. We compare the performance of the policies trained 1) with only dropout during training ($p = 0.001$), 2) with the dropout during training ($p = 0.001$) plus dropout during rollout ($p = 0.0001$), 3) with the policy trained with Roll-Drop only. In Figure 7(a), we can observe that the introduction of dropout during training, even in conjunction with dropout during rollouts is detrimental.

6.2. Dropout probability and convergence

The classical usage of dropout in supervised learning is to regularise the learning of the employed networks Srivastava et al. (2014). Randomly dropping units from the Neural Network (NN) during

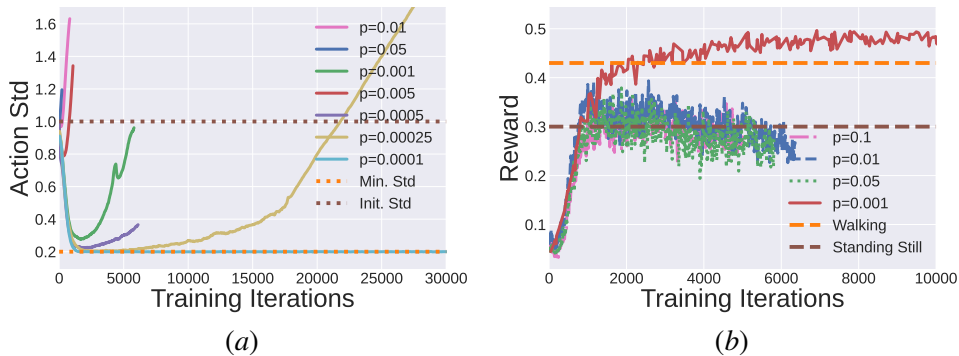


Figure 3: In Figure 3(a) the introduction of dropout during rollouts heavily affects the convergence, in this figure we show the effect of different dropout probabilities on the standard deviation used to sample actions in policy gradient algorithms. In Figure 5(c) alongside dropout during rollouts, we investigate the effects of dropout during training. When the dropout probability is too high the policy displays a standing-still behaviour, which suggests its inability to correlate inputs to outputs.

training prevents them from co-adapting, thus significantly reducing overfitting; at test time, then, the dropout is removed to approximate averaging the predictions of all these partial networks by using a complete network with smaller weights. A second dropout application is to approximate a Bayesian network Gal and Ghahramani (2016): applying dropout at inference time can generate multiple predictions from the same input through multiple inferences. This gives us a probability distribution of the outputs which we can then analyse. For such applications dropout probability typically varies between 20% and 50%, but can reach values up to 80% Srivastava et al. (2014).

As can be seen from Figure 3(a), in the case of Roll-Drop the probability is much lower: 0.01%. In fact, unlike supervised and semi-supervised learning, RL does not provide any target and the policy loss depends entirely on the actions taken by the policy itself. The better states the policy explores, the higher the reward it will receive; conversely exploring bad states can result in the policy exploring a wider actions space and eventually catastrophically diverging to even worse states. Moreover, when some neurons are dropped during rollouts the noise introduced affects the following state of the episode. Considering the latter in conjunction with having a moving target, it is clear that RL is more sensitive to dropout probabilities and that lower dropout probabilities are expected.

Figure 3(a) shows how the training diverges for dropout probabilities $\in [0.00025, 0.01]$, while for $p = 0.0001$ it converges to a stable behaviour. In fact, policy gradient algorithms like PPO (Schulman et al., 2017) explore the action space sampling from a distribution $\mathcal{N}(\mu, \sigma)$, where μ is the action a output of π , while σ is a learnt parameter. Figure 3(a) shows the σ for policies trained with different levels of dropout during rollouts (same random seed) – the initial $\sigma_0 = 1.$, and it is capped to $\sigma_{min} = 0.2$; high probabilities of dropout are responsible for the divergence.

6.3. How is Roll-Drop affecting the training?

Figure 1 depicts the effects of Roll-Drop on the training: the policies with and without dropout observe the same initial state (same random seed), and produce the same action until the first dropout is triggered. After this event, the two trainings take different trajectories \mathcal{T} and \mathcal{T}' , as the policies output different actions and the robots experience different states. We investigated




Figure 4: Policies trained with *No randomisation* (left) and Roll-Drop (right), both trained on flat ground in simulation with $K_p = 20$, and tested on the hardware with $K_p = 15$. This gain change highlights a case of system uncertainty, demonstrating that Roll-Drop can address sim-to-real gaps.

this further in Figure 5 by recording states and actions for the first 3000 training iterations (adopting 128 parallel environments and episodes of 4 [s]) for Roll-Drop with associated probabilities $\in [0., 0.002, 0.0001]$. We show these distribution shifts for some states and actions: the joint position of HR_KFE in Figure 5(a), the joint velocity of HL_KFE in Figure 5(b), and the action of HR_HAA in Figure 5(c). These histograms demonstrate how such a tiny dropout probability – when compared to supervised/semi-supervised/non-supervised learning– affects the training: For $p = 0.0001$ the distributions of states and actions are different from $p = 0.$, while for $p = 0.002$ the training is clearly diverging (Figure 3(a)) with most of the states and actions distributed close to the joint position and velocity limits. Further evidences of the divergence are provided in Figures 5(d) to 5(f), where we show the mean across all environments and time-steps for each of the first 3000 iterations.

6.4. Different random seeds

We investigated how consistent the training is when different random seeds are used: We trained five policies without any randomisation and five policies with Roll-Drop across five different seeds, and compared the total reward of both groups in terms of mean and standard deviation. We considered the *No Randomisation* setting as the perfect candidate for this comparison since it is massively overfitting to the simulation environment. From Figure 7(b), we discovered that Roll-Drop (blue line) is on average performing better, because of its higher robustness across different seeds, and this is also supported by the smaller standard deviation, when compared to *No Randomisation*. However, as expected, *No Randomisation* is in absolute value performing better than Roll-Drop, but only for the seed the environment was originally tuned on; while for other seeds it gained lower rewards, and it has a more spread standard deviation.

6.5. Training and deployment mismatch

Apart from increasing the robustness to observation noise, Roll-Drop is also providing the policies with additional flexibility to external perturbations. Indeed, we trained two more policies -with and without Roll-Drop- with $K_p = 20$ and we deployed them on the hardware using $K_p = 15$, the target velocity command is $s_c = [0., 0., 0.]$. As can be seen from Figure 4, the policy trained without any sort of randomisation is not able to stand, while on the other hand Roll-Drop allows the policy to find equilibrium and to better follow the velocity command. Quantitative advantages of adopting Roll-Drop for the experiments above are provided in Figure 6, where we show better velocity command tracking (both linear and angular), and lower joint velocity usage.

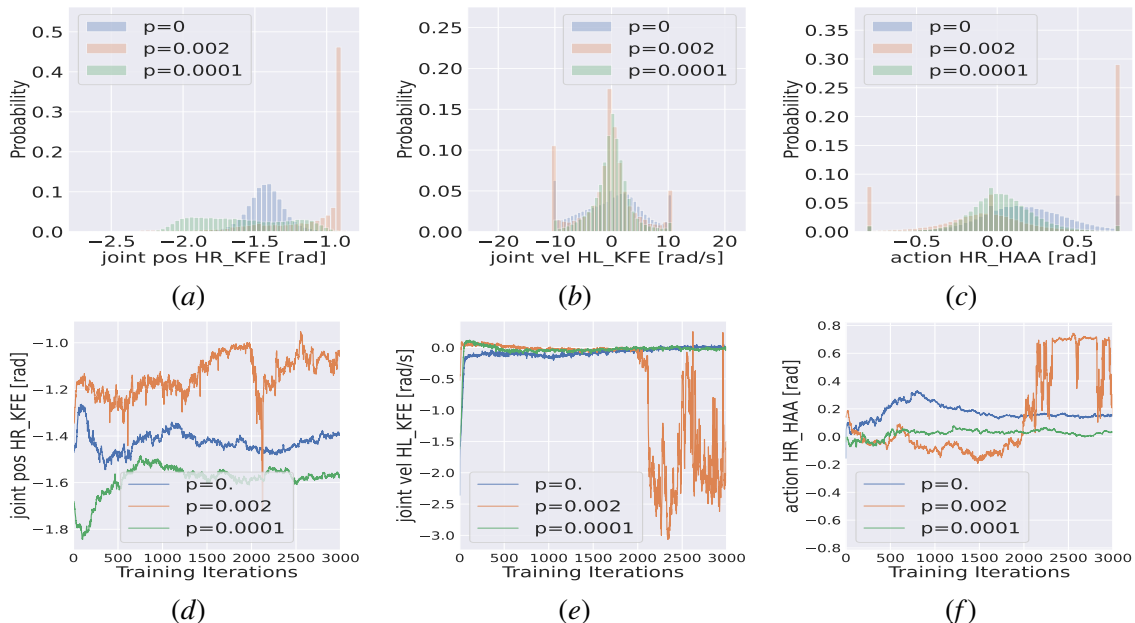


Figure 5: (a) to (c) show some state/action distributions for Roll-Drop with three different probabilities: 0 (no Roll-Drop), 0.002, and 0.0001. When $p = 0$ the policy is over-fitting to a fixed simulation environment, when $p = 0.0001$ the policy converges to the desired behaviour, but exploring a different state/action space compared to $p = 0$, and finally, when $p = 0.002$ it diverges catastrophically and it explores bad portions of the state/action spaces (often the joint position/velocity limits) as can be seen from the orange histograms. Evidences of the divergence are provided in (d) to (f), which represent the mean for each state/action along the initial 3000 iterations considered.

7. Conclusion

In this work we show how to account for observation noise without tuning randomisation distributions for each of the states/sensors as is commonly used in DRL. This can be simply done by including dropout during rollouts (Roll-Drop) in the network architecture and by tuning a single parameter: the dropout probability. In fact, by turning on and off neurons during the rollouts we show a considerable improvements in noise-injection robustness (200%), and a success rate of 80% when 25% noise is injected. Alongside the results we present a thorough analysis to explain the effects of different dropout implementations and associated probabilities on performances, convergence, and state/action distributions. The approach was also validated on the hardware and tested on board of the Unitree A1 quadruped robot. Although we conducted a thorough ablation study, future research will explore the impact of changing the position of the dropout layer in the network on the controller’s overall performance.

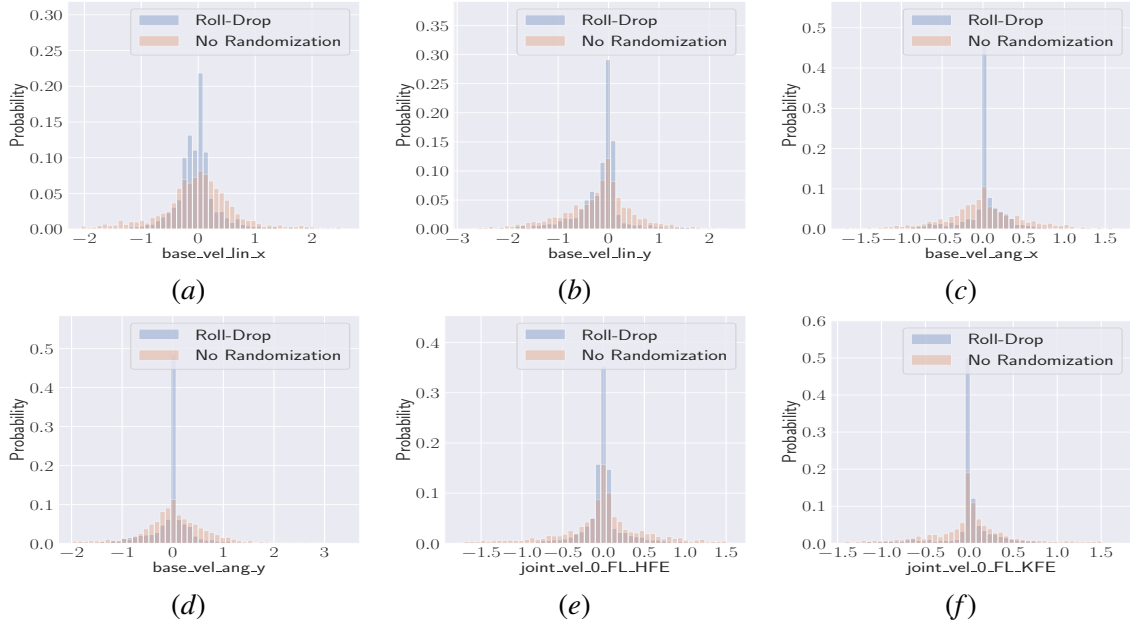


Figure 6: These policies are trained in simulation with $K_p = 20$ and deployed on the hardware (Uni-tree A1) using $K_p = 15$. In this context Roll-Drop demonstrated better performances in tracking the velocity command and in expending less velocity at the joints.




Figure 7: (a) Comparing the success rate of different methods in the presence of noise in the observations. Roll-Drop performs more than twice as good as other methods, retaining 80% success rate with more than 25% injected noise. (b) we trained the Roll-Drop and the *No Randomisation* policies using 5 different random seeds, and here we compare their total rewards. Roll-Drop is more consistent, demonstrating on average higher total reward and a smaller standard deviation. Conversely, *No Randomisation* is more sensitive to the random seed used: highest total reward on the seed on which the environment was tuned on, but the performance degraded when other seeds were adopted.

References

- Steven Bohez, Saran Tunyasuvunakool, Philemon Brakel, Fereshteh Sadeghi, Leonard Hasenclever, Yuval Tassa, Emilio Parisotto, Jan Humplik, Tuomas Haarnoja, Roland Hafner, Markus Wulfmeier, Michael Neunert, Ben Moran, Noah Siegel, Andrea Huber, Francesco Romano, Nathan Batchelor, Federico Casarini, Josh Merel, Raia Hadsell, and Nicolas Heess. Imitate and repurpose: Learning reusable robot movement skills from human and animal behaviors, 2022. URL <https://arxiv.org/abs/2203.17138>.
- Luigi Campanaro, Siddhant Gangapurwala, Wolfgang Merkt, and Ioannis Havoutis. Learning and deploying robust locomotion policies with minimal dynamics randomization, 2022. URL <https://arxiv.org/abs/2209.12878>.
- Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR, 2016.
- Siddhant Gangapurwala, Mathieu Geisert, Romeo Orsolino, Maurice Fallon, and Ioannis Havoutis. Rloc: Terrain-aware legged locomotion using reinforcement learning and optimal control. *IEEE Transactions on Robotics*, 38(5):2908–2927, 2022. doi: 10.1109/TRO.2022.3172469.
- Geoffrey E. Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *CoRR*, abs/1207.0580, 2012. URL <http://arxiv.org/abs/1207.0580>.
- Christian Hubicki, Jesse Grimes, Mikhail Jones, Daniel Renjewski, Alexander Spröwitz, Andy Abate, and Jonathan Hurst. Atrias: Design and validation of a tether-free 3d-capable spring-mass bipedal robot. *The International Journal of Robotics Research*, 35(12):1497–1521, 2016. doi: 10.1177/0278364916648388. URL <https://doi.org/10.1177/0278364916648388>.
- Marco Hutter, Christian Gehring, Dominic Jud, Andreas Lauber, C. Dario Bellicoso, Vassilios Tsounis, Jemin Hwangbo, Karen Bodie, Peter Fankhauser, Michael Bloesch, Remo Diethelm, Samuel Bachmann, Amir Melzer, and Mark Hoepflinger. Anymal - a highly mobile and dynamic quadrupedal robot. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 38–44, 2016. doi: 10.1109/IROS.2016.7758092.
- Jemin Hwangbo, Joonho Lee, and Marco Hutter. Per-contact iteration method for solving contact dynamics. *IEEE Robotics and Automation Letters*, 3(2):895–902, 2018. URL www.raisim.com.
- Jemin Hwangbo, Joonho Lee, Alexey Dosovitskiy, Dario Bellicoso, Vassilios Tsounis, Vladlen Koltun, and Marco Hutter. Learning agile and dynamic motor skills for legged robots. *Science Robotics*, 4(26):eaau5872, 2019. doi: 10.1126/scirobotics.aau5872. URL <https://www.science.org/doi/abs/10.1126/scirobotics.aau5872>.
- Nick Jakobi, Phil Husbands, and Inman Harvey. Noise and the reality gap: The use of simulation in evolutionary robotics. In Federico Morán, Alvaro Moreno, Juan Julián Merelo, and Pablo Chacón, editors, *Advances in Artificial Life*, pages 704–720, Berlin, Heidelberg, 1995. Springer Berlin Heidelberg. ISBN 978-3-540-49286-3.

- Gwanghyeon Ji, Juhyeok Mun, Hyeongjun Kim, and Jemin Hwangbo. Concurrent training of a control policy and a state estimator for dynamic and robust legged locomotion. *IEEE Robotics and Automation Letters*, 7(2):4630–4637, 2022. doi: 10.1109/LRA.2022.3151396.
- Ashish Kumar, Zipeng Fu, Deepak Pathak, and Jitendra Malik. RMA: Rapid motor adaptation for legged robots. In *Robotics: Science and Systems*, 2021.
- Joonho Lee, Jemin Hwangbo, and Marco Hutter. Robust recovery controller for a quadrupedal robot using deep reinforcement learning. *CoRR*, abs/1901.07517, 2019. URL <http://arxiv.org/abs/1901.07517>.
- Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning quadrupedal locomotion over challenging terrain. *Science Robotics*, 5(47):eabc5986, 2020. doi: 10.1126/scirobotics.abc5986. URL <https://www.science.org/doi/abs/10.1126/scirobotics.abc5986>.
- Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, and Gavriel State. Isaac gym: High performance GPU based physics simulation for robot learning. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2021. URL https://openreview.net/forum?id=fgFBtYgJQX_.
- Matias Mattamala, Nived Chebrolu, and Maurice Fallon. An efficient locally reactive controller for safe navigation in visual teach and repeat missions. *IEEE Robotics and Automation Letters*, 7(2): 2353–2360, 2022. doi: 10.1109/LRA.2022.3143196.
- Takahiro Miki, Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning robust perceptive locomotion for quadrupedal robots in the wild. *Science Robotics*, 7(62):eabk2822, 2022. doi: 10.1126/scirobotics.abk2822. URL <https://www.science.org/doi/abs/10.1126/scirobotics.abk2822>.
- Xue Bin Peng and Michiel van de Panne. Learning locomotion skills using deeprl: Does the choice of action space matter? In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 1–13, 2017.
- Nikita Rudin, David Hoeller, Philipp Reist, and Marco Hutter. Learning to walk in minutes using massively parallel deep reinforcement learning. In Aleksandra Faust, David Hsu, and Gerhard Neumann, editors, *Proceedings of the 5th Conference on Robot Learning*, volume 164 of *Proceedings of Machine Learning Research*, pages 91–100. PMLR, 08–11 Nov 2022. URL <https://proceedings.mlr.press/v164/rudin22a.html>.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *CoRR*, abs/1707.06347, 2017. URL <http://arxiv.org/abs/1707.06347>.
- Jonah Siekmann, Yesh Godse, Alan Fern, and Jonathan Hurst. Sim-to-real learning of all common bipedal gaits via periodic reward composition. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7309–7315, 2021. doi: 10.1109/ICRA48506.2021.9561814.

- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15(1): 1929–1958, jan 2014. ISSN 1532-4435.
- Jie Tan, Tingnan Zhang, Erwin Coumans, Atil Iscen, Yunfei Bai, Danijar Hafner, Steven Bohez, and Vincent Vanhoucke. Sim-to-real: Learning agile locomotion for quadruped robots. In *Proceedings of Robotics: Science and Systems*, Pittsburgh, Pennsylvania, June 2018. doi: 10.15607/RSS.2018.XIV.010.
- Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 23–30, 2017. doi: 10.1109/IROS.2017.8202133.
- Eugene Valassakis, Zihan Ding, and Edward Johns. Crossing the gap: A deep dive into zero-shot sim-to-real transfer for dynamics. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5372–5379, 2020. doi: 10.1109/IROS45743.2020.9341617.
- Chuanyu Yang, Kai Yuan, Qiuguo Zhu, Wanming Yu, and Zhibin Li. Multi-expert learning of adaptive legged locomotion. *Science Robotics*, 5(49):eabb2174, 2020.
- Fangzhou Yu, Ryan Batke, Jeremy Dao, Jonathan Hurst, Kevin Green, and Alan Fern. Dynamic bipedal maneuvers through sim-to-real reinforcement learning, 2022. URL <https://arxiv.org/abs/2207.07835>.