

Learning Torque Control for Quadrupedal Locomotion

Shuxiao Chen^{*1}, Bike Zhang^{*1,2}, Mark W. Mueller¹, Akshara Rai² and Koushil Sreenath¹

Abstract—Reinforcement learning (RL) has become a promising approach to developing controllers for quadrupedal robots. Conventionally, an RL design for locomotion follows a position-based paradigm, wherein an RL policy outputs target joint positions at a low frequency that are then tracked by a high-frequency proportional-derivative (PD) controller to produce joint torques. In contrast, for the model-based control of quadrupedal locomotion, there has been a paradigm shift from position-based control to torque-based control. In light of the recent advances in model-based control, we explore an alternative to the position-based RL paradigm, by introducing a torque-based RL framework, where an RL policy directly predicts joint torques at a high frequency, thus circumventing the use of a PD controller. The proposed learning torque control framework is validated with extensive experiments, in which a quadruped is capable of traversing various terrain and resisting external disturbances while following user-specified commands. Furthermore, compared to learning position control, learning torque control demonstrates the potential to achieve a higher reward and is more robust to significant external disturbances. To our knowledge, this is the first sim-to-real attempt for end-to-end learning torque control of quadrupedal locomotion.

I. INTRODUCTION

Legged locomotion is a longstanding challenge in the field of robotics. While the high degree of freedom, underactuation, and hybrid nonlinear dynamics of legged robots are particularly difficult, they have the advantage of being able to traverse rough terrain, making them highly useful in certain applications. Among the different types of legged robots, quadrupeds have received considerable attention over the last few decades and have achieved significant advancements.

Since the 1980s, model-based control has been a primary method for achieving quadrupedal locomotion [26]. This approach uses a dynamics model to develop advanced controllers that enable a quadruped to perform a range of locomotion behaviors with rigorous stability and safety guarantees. Model-based control for quadrupeds has undergone a significant transformation, shifting from position control [25] to torque control [7]. This change has led to more dynamic and robust locomotion behaviors, indicating the superiority of torque control. On the other hand, learning-based approaches have recently emerged as a promising tool for quadrupedal locomotion, offering great potential for navigating rough terrain [10], [16], [20]. A notable difference between the learning-based approach and model-based control is that the former typically uses joint positions



Fig. 1: A quadrupedal robot walks on a hill with a torque-based RL policy. The experimental video can be found at <https://youtu.be/18v3Z9BD-W8>.

as the action space, whereas the latter favors torque control. In addition, it is worth noting that learning torque control has been utilized in various robotic applications beyond locomotion, e.g., manipulators [37] and quadrotors [41], with promising results. Inspired by the successful transition of model-based control, we explore an alternative approach for learning-based quadrupedal locomotion, where an RL policy directly predicts joint torques, instead of joint positions with a low-level PD controller.

A. Related Work

Learning quadrupedal locomotion has been widely studied in the past several decades. In the DARPA learning locomotion program, a hierarchical apprenticeship learning method was developed for quadrupedal walking [13]. More recently, deep reinforcement learning has emerged as a viable approach for quadrupedal locomotion. For instance, in [36], a walking policy was trained in simulation and successfully transferred to a real quadrupedal robot. Moreover, imitation from animals was explored to produce learned quadrupedal motions in [22]. Several studies have also explored the use of model-based control in learned quadrupedal locomotion, as evidenced by works such as [6], [17], [8]. In addition to learning only in simulation, a substantial part of learning can also happen in the real world [33], [38]. Furthermore, a series of learning quadrupedal locomotion works have achieved state-of-the-art performance for traversing challenging terrain [10], [16], [20]. Additionally, other works have demonstrated different learned locomotion behaviors [31], [3], [4], [40], [5] and explored various learning approaches [35], [12], [14].

^{*} These authors contributed equally to this work.

¹ These authors are with the Department of Mechanical Engineering, University of California, Berkeley, CA, 94720, USA, {shuxiao.chen, bikezhang, mwm, koushils}@berkeley.edu

² These authors are with Meta AI, CA, 94025, USA, {bikezhang, akshararai}@meta.com

For the reinforcement learning-based legged locomotion, there are two main approaches in terms of the action space design: one is joint position [36], [16] and the other is joint torque [18], [29]. When joint positions are used as actions, the learned policy outputs target joint angles which are then sent to a PD controller. The PD controller is treated as a low-level tracking module, and the PD gains are usually tuned manually to track desired positions well. In addition, a prevalent variation is the residual joint position, where joint positions consist of heuristic or optimization-based joint positions and learned residual joint positions [16]. When the action space is joint torque, actions are directly applied to motors to actuate a robot. However, this action choice has only been demonstrated in simulation.

There has been some research to investigate the efficacy of both action choices. A study on different proportional gains for the PD controller was conducted in [39], which suggests that high proportional gains can cause instability on the robot while low proportional gains result in a policy behaving like a torque controller that leads to a more robust policy. Nevertheless, low proportional gains have larger tracking errors. In addition, the derivative gain for the PD controller was studied in [34], which validates that a small derivative gain leads to learning instabilities and a large derivative gain prevents the robot from tracking a target velocity. Thus, in order to have adaptive stiffness and damping at the joints, the PD gains can also be learned as part of the action space, as demonstrated in [32]. However, switching between different PD gains can potentially produce nonsmooth joint torques and cause instabilities of the robot.

In spite of the exploration in the design of PD controller, learning torque control was studied in [23] in a simulation setting, where a comparison among different action spaces was conducted, including joint positions, joint torques, joint velocities, and activations for musculotendon units. Note that in [23, Fig. 4], joint torques are compared with joint positions at the same frequency for a quadruped. However, joint torques need to be executed at a higher frequency in order to achieve comparable performance to joint positions. Moreover, this figure shows that the performance of joint torques becomes better when the query rate increases, which implies the potential of learning torque control. Furthermore, in [27, Sec. 8], a simulation study of joint positions and joint torques indicates that, while learning joint position control leads to a quicker early learning rate, learning joint position control may converge to lower-reward behaviors and may hinder the ability to adapt to rapid changes.

Building upon the previous studies, in this work, we investigate the possibility of joint torques as the action space for a reinforcement learning policy with application to quadrupedal locomotion. To the best of our knowledge, torque control for real-world quadrupedal locomotion using end-to-end deep reinforcement learning has not been demonstrated yet, and the potential of learning torque control for quadrupedal locomotion is still underexplored.

B. Contributions

The contributions of this work with respect to prior literature are summarized as follows:

- We present a sim-to-real learning torque control approach for quadrupedal locomotion, where an RL policy streams joint torques directly, without a PD controller.
- We validate the proposed approach through extensive experiments, in which a quadruped is capable of traversing different terrain and resisting external disturbances.
- We illustrate that learning torque control has the potential to achieve higher rewards, compared to the conventional learning position control.
- We demonstrate that a quadruped trained with torque-based RL is capable of resisting larger disturbances when compared with the position-based approach.
- We further discuss the key features of learning torque control and provide implementation suggestions.

II. LEARNING TORQUE CONTROL FOR LOCOMOTION

In this section, we present our end-to-end learning torque control framework. Fig. 2 illustrates the observations and policy architecture used to predict actions. Compared with the common position-based method, our proposed torque-based policy is queried at a high frequency, and it directly outputs joint torques, avoiding the use of a PD controller.

A. Simulation Environment

We choose Isaac Gym [19], [28] as our simulator. We train our policy end-to-end with 4096 parallel environments simulated on an NVIDIA GeForce RTX 3090 GPU. In addition, we use a curriculum of rough terrain in the simulation to improve the robustness of our policy.

B. Problem Formulation

We formulate quadrupedal locomotion as a reinforcement learning problem where the objective is to learn a policy π that allows an agent to maximize its expected return for a given task. At each time step t , a state s_t is observed by the agent from the environment, and an action $\mathbf{a}_t \sim \pi(\mathbf{a}_t | s_t)$ is sampled from the policy π . The agent then applies this action and drives its current state to a new state s_{t+1} and a scalar reward $r_t = r(s_t, \mathbf{a}_t, s_{t+1})$ is received. This process is iterated to generate a trajectory $\tau = \{(s_0, \mathbf{a}_0, r_0), (s_1, \mathbf{a}_1, r_1), \dots\}$. Therefore, the goal is to learn a policy that maximizes the agent's expected return $J(\pi)$ defined as

$$J(\pi) = \mathbb{E}_{\tau \sim p(\tau | \pi)} \left[\sum_{t=0}^{T-1} \gamma^t r_t \right],$$

where T is the time horizon for each episode, and $\gamma \in [0, 1]$ is a discount factor. $p(\tau | \pi)$ represents the likelihood of a trajectory τ under a policy π ,

$$p(\tau | \pi) = p(s_0) \prod_{t=0}^{T-1} p(s_{t+1} | s_t, \mathbf{a}_t) \pi(\mathbf{a}_t | s_t),$$

where $p(s_0)$ is the initial state distribution, and $p(s_{t+1} | s_t, \mathbf{a}_t)$ denotes the state transition.




Fig. 2: The proposed learning torque control framework. The common position-based method (top left) has two main modules where a neural network (blue) runs at a low frequency and a PD controller (green) executes at a high frequency. The proposed torque-based method (top right) merges the two modules in the position-based method into a high-frequency neural network module (red), as the RL policy is queried at a high frequency to directly output torques. A detailed illustration of the high-frequency neural network module is shown at the bottom, consisting of training in the simulation environment and test on the real quadruped. During the training, we add noise and disturbance to facilitate sim-to-real transfer. During the test, the state information is collected through proprioceptive sensors with appropriate preprocessing. The comprehensive explanation of the framework is presented in Sec. II.

C. Observation Space

We choose a concise representation of the robot that can prove the advantage of learning direct torque control, alleviate overfitting to simulated dynamics, and facilitate sim-to-real transfer [36], excluding any history of proprioceptive sensor measurements or privileged information. The observation space consists of 7 parts and they together form a 48×1 vector. The different components consist of robot base linear and angular velocities, orientation, joint positions and velocities, last action, and user commands. Comprehensive definitions of the observation space can be found in Table I. Base linear and angular velocities and projected gravity are in the robot frame. The projected gravity is a unit vector capturing the robot's orientation with respect to gravity. Desired forward velocity, lateral velocity, and yaw rate are sent to the controller as the user commands to modify the policy behavior online. Joint position, joint velocity, and last action are each 12-dimensional corresponding to the hip, thigh, and calf joints of each of the four legs of the quadruped. The last action is the last neural network output without being multiplied by the action scale.

D. Action Space

The action space of our model is a 12×1 vector consisting of unscaled joint torques. This vector is then multiplied by

TABLE I: Detailed definitions of the components forming the observation space. The observation space is divided into 7 observation items whose corresponding symbols are listed side by side. The scale factors stand for the coefficients that are multiplied to the observation items as a preprocessing step. The upper bounds of the zero-mean uniform noise added to the observation are also listed below.

Observation item	Symbol	Scale factor	Noise
Base linear velocity	$\mathbf{v}_{\text{base}} = [v_x \ v_y \ v_z]^T \in \mathbb{R}^3$	2.0	0.2
Base angular velocity	$\boldsymbol{\omega}_{\text{base}} = [\omega_x \ \omega_y \ \omega_z]^T \in \mathbb{R}^3$	0.25	0.2
Projected gravity	$\mathbf{g}_{\text{proj}} = [g_x \ g_y \ g_z]^T \in \mathbb{R}^3$	1.0	0.05
Joint position	$\mathbf{q} = [q_1 \ \dots \ q_{12}]^T \in \mathbb{R}^{12}$	1.0	0.01
Joint velocity	$\dot{\mathbf{q}} = [\dot{q}_1 \ \dots \ \dot{q}_{12}]^T \in \mathbb{R}^{12}$	0.05	1.5
User command	$\mathbf{u} = [v_x \ v_y \ \omega_z]^T \in \mathbb{R}^3$	$[2.0 \ 2.0 \ 0.25]^T$	0
Last action	$\mathbf{a}_{\text{last}} = [\tau_1 \ \dots \ \tau_{12}]^T \in \mathbb{R}^{12}$	1.0	0

a constant action scale to produce the actual desired torque values sent to the 12 motors.

The intuition behind the action scaling is as follows: typical position-based RL policies usually add or subtract position values from a reference or nominal position vector, like a standing pose for the robot. There are two reasons for this choice. First, learning a residual on a nominal configuration of joint positions makes it easier for the agent to survive during the early stages of training by giving it a

TABLE II: Detailed definitions of the proposed reward terms. Variables marked with an asterisk stand for target values. Δt is defined by the inverse of the control frequency, in this case, 0.005 s. Positive weights mean encouragement to such behaviors and negative ones penalty. In the hip (thigh) reward, H (T) is the set of 4 hip (thigh) motor indices.

Reward term	Expression	Weight
Base linear velocity x - y	$\exp(-\ \mathbf{v}_{\text{base},xy}^* - \mathbf{v}_{\text{base},xy}\ ^2/0.25)$	$1.0\Delta t$
Base linear velocity z	$-v_{\text{base},z}^2$	$-2.0\Delta t$
Base angular velocity x - y	$-\ \boldsymbol{\omega}_{\text{base},xy}\ ^2$	$-0.05\Delta t$
Base angular velocity z	$\exp(-(\omega_{\text{base},z}^* - \omega_{\text{base},z})^2/0.25)$	$0.5\Delta t$
Orientation	$\ \mathbf{g}_{\text{projected},xy}\ ^2$	$-1.0\Delta t$
Joint torque	$\ \boldsymbol{\tau}\ ^2$	$-0.0002\Delta t$
Joint acceleration	$\ \dot{\mathbf{q}}_{\text{last}} - \dot{\mathbf{q}}\ ^2$	$-0.0000025\Delta t$
Feet air time	$\sum_{i=1}^4 (t_{\text{swing},i} - 0.5)$	$2.0\Delta t$
Collision	$\sum_{i=1}^4 c_{\text{knee},i} \text{ for } c_{\text{knee},i} \in \{0, 1\}$	$-1.0\Delta t$
Action rate	$\ \mathbf{a}_{\text{last}} - \mathbf{a}\ ^2$	$-0.01\Delta t$
Hip	$\sum_{h \in H} q_h^* - q_h $	$-1.0\Delta t$
Thigh	$\sum_{h \in H} q_t^* - q_t $	$-1.0\Delta t$

prior standing pose. Secondly, it also helps establish a good starting point for policy exploration because desired gaits often occur around this nominal configuration. However, determining nominal torques is difficult for torque control, since these might vary with different environments and initial configurations. Hence, we remove the use of reference or nominal torques and instead use a scaling vector to give final torques to the motor controller.

Due to the fact that the numerical range of torque control (e.g., -33 to 33 Nm being the safe joint torque range on the A1 robot) is much larger than that of position control (e.g., -3.14 to 3.14 rad), we use a large value of 10 as the action scale. This action scale strikes a good balance between exploration and exploitation in the absence of nominal torques, and it also helps stabilize the robot in the early stage of training. Eventually, all final torques are clamped to the range $[-33, 33]$ Nm in both simulation and real-world tests to ensure the safety of the robot whilst retaining enough leeway for agile locomotion.

E. Rewards

Similar to having no nominal configurations for the action space, no reference trajectories or demonstrations are used in our reward functions either. The learned policy is able to evolve spontaneously from high-level task rewards, and generates gaits that are suitable for high-frequency torque control without any additional constraints to reference motions.

Overall, the total reward function encourages the agent to follow a given set of high-level user commands, whilst satisfying additional behavioral, heuristic, and cosmetic requirements. Detailed definitions of the individual reward terms can be seen in Table II. The hip and thigh reward terms drive the positions of the joints close to nominal positions. This creates “natural looking” gaits on the robot, enhancing the sim-to-real transfer probability of learned policies.

F. Policy Architecture and Training

The policy is modeled with a multilayer perceptron (MLP) which has 3 layers of 512, 256, and 128 exponential linear units (ELUs). The standard deviation of the action distribution is set to a fixed value. The value function is modeled by a separate network with the same configuration as the policy network. The policy is optimized with proximal policy optimization (PPO) [30]. Note that we do not choose a complicated architecture, such as a teacher-student structure in [16], since we would like to use a simple architecture to validate our learning torque control framework.

G. Domain Adaptation

Randomization of ground frictions and external disturbances is included to facilitate the building of robustness. A friction coefficient is sampled uniformly in the range $[0.5, 1.25]$ and assigned to each simulated robot. External pushes are exerted on each robot every 15 s by sampling and setting the instantaneous x - y velocity of the robot base from $[-1.0, 1.0]$ m/s. Note that uncertainties of the physical parameters, e.g., dynamics randomization [2], [21], can be viewed as external forces applied to the robot [24], [36]. The linear velocity commands are also randomized in range $[-1.0, 1.0]$ m/s in both x - and y -directions, and the yaw rate is randomized in range $[-3.14, 3.14]$ rad/s, in order to produce a goal-conditioned RL policy.

III. EXPERIMENTAL SETUP AND RESULTS

Having presented the end-to-end learning torque control framework, we now evaluate it with various experiments. We first introduce the experimental apparatus and a baseline policy. Then, real-world experiments and comparisons are demonstrated. Finally, we provide an in-depth discussion on the learning torque control framework.

A. Apparatus

We use a Unitree A1 quadrupedal robot to evaluate our end-to-end learning torque control framework. The A1 robot has 18 degrees of freedom (DoFs) with 12 actuated joints and 6 generalized coordinates for the floating base. There is an NVIDIA Jetson TX2 inside the robot body to offer extra computing power. The TX2 has a 256-core GPU for neural network inference implemented with Mobile Neural Network (MNN) [11], an open-source inference engine that is lightweight and supports various hardware backends. In addition, our policy runs completely onboard.

B. Baseline

We develop a position-based RL policy for the same quadruped as a baseline. This policy uses joint positions as the action space and runs at 50 Hz, followed by a low-level PD controller tracking at 200 Hz, which is the same frequency as the torque-based RL policy. For the PD controller, the proportional gain is 20, and the derivative gain is 0.5. This position-based RL policy is trained with the same level of noise and disturbance as well as the terrain curriculum as the torque-based RL policy. In order to make a fair comparison,



Fig. 3: The proposed end-to-end learning torque control framework allows the A1 robot to traverse various terrains, such as (a) inclined soft lawn, (b) pebble path, (c) compliant terrain, and resist external disturbances such as (d) recovering from a kick. All the tests use the same RL policy queried at 200 Hz.

we make the reward function of the position-based RL close to the torque-based RL one. However, the same reward function does not work directly for position-based RL policy. We modify the reward function in a minimal way which excludes the orientation, hip, and thigh reward terms, and changes the weight of the feet air time reward from 2.0 to 1.0. The action scale for this position-based policy is 0.25 with a nominal position. This position-based policy also allows a quadruped robot to traverse different terrains.

C. Results

We conduct experiments both in simulation and on hardware to extensively study our proposed learning torque control framework. In the simulation, we compare the difference between learning torque control and learning position control in terms of their learning curves. On hardware, we validate the learned torque control policy on various terrains with user-specified commands. The snapshots of the learned locomotion behaviors are shown in Fig. 1 and Fig. 3.

In Fig. 3a, a quadruped robot walks on an inclined soft lawn on campus, which shows that the learned torque policy allows the quadruped to locomote over non-flat terrain while following commands. In Fig. 3b, the quadruped robot walks on a pebble path, where the pebbles are not embedded in the ground and can be slippery. While the robot walks forward, the pebbles collide with the feet, which causes additional disturbances to the robot. In this experiment, the robot gets stuck in pebbles for a moment. The operator increases the forward velocity command, and then the robot is able to get out of the pebbles and keeps walking forward. Fig. 3c demonstrates that our torque policy enables the quadruped robot to walk over a compliant terrain, i.e., an air mattress. A similar test was conducted in [5] using a position-based RL policy with an adaptive control architecture. In addition to walking on various terrains, we also test the robot with external disturbances as shown in Fig. 3d. We kick the robot from different angles when it is standing or walking, and it is capable of resisting a certain amount of external forces. All of these locomotion behaviors are performed with the same torque-based RL locomotion policy queried at 200 Hz.

In order to study the difference between learning position control and learning torque control, we carry out real-world experiments to compare their action tracking performances.




Fig. 4: **Action tracking comparison** between learning position control (top) and learning torque control (bottom) on a real quadruped robot for the rear left calf joint. The green lines depict the target values queried from RL policies with post-processing, e.g., scales. The orange lines illustrate the actual values measured from the real quadruped robot.

We perform real-world experiments on the same A1 robot given the same forward velocity command at 0.5 m/s. The action tracking result for the rear left calf joint is shown in Fig. 4, where the top plot illustrates the result of learning position control and the bottom plot demonstrates the result of learning torque control. The green lines denote the target values produced from RL policies with post-processing, and the orange lines indicate the actual values measured from the hardware. Learning position control has a larger tracking error, whereas learning torque control has a much better tracking performance. It is worth noting that this comparison does not indicate that learning position control has worse performance, since the position-based RL can adjust the target position so that the actual position achieves the goal during training. Here, we only aim to demonstrate a major difference between the two methods.

We were concerned that joint torques need to be changed rapidly between the stance phase and the swing phase for a quadruped robot, compared to joint positions. Thus, it may pose challenges to RL to learn to control a quadruped robot




Fig. 5: **Learning curve comparison** between learning position control and learning torque control. We choose three different proportional gains, i.e., 40, 60, and 80, for the position-based RL, which are denoted as orange, green, and red, respectively, and the torque-based RL is depicted as blue.

using joint torques as the action space. Fig. 4 indicates that RL is powerful enough to handle this rapid change in the action space. In addition, we were concerned that the action of RL might not be smooth enough to be fed to motors. For learning position control, there is a low-level PD controller to track the action of RL, and it can locally smooth out jerky values. Fig. 4 demonstrates that torque-based RL produces smooth values for motors. Although we occasionally notice some jerky values, the motors on the hardware are generally tolerant. There is no need to add another layer, e.g., a PD controller or a low pass filter, to smooth out the RL actions.

Furthermore, we compare the normalized learning curves for learning position control and learning torque control, as shown in Fig. 5. We simplify the reward function design for this comparison, which only consists of the base velocity tracking reward. We choose three different proportional gains, i.e., 40, 60, and 80, which are illustrated as orange, green, and red lines, and the derivative gain is 0.5. The torque-based RL is depicted as the blue line. The learning curves for the position-based RL change with respect to different proportional gains, which indicates that proportional gains have an impact on the learning process. In contrast, while torque-based RL shows a slow learning rate at the early stage, it converges to a higher return, which shows that torque-based RL has the potential to achieve a higher expected return, compared to position-based RL. This result coincides with the findings presented in [27], which conducts a similar comparison but using a much simpler system.

In addition, we compare the robustness between the position-based RL policy and the torque-based RL policy. We perform a real-world experiment where the robot stands while getting consecutive kicks, and the result is presented in Fig. 6. The green, blue, and red lines show the measured, nominal, and minimal base heights, respectively. The robot

base heights are calculated by averaging the z -direction foot positions in the robot frame obtained by forward kinematics (FK). Since the robot body has non-negligible thickness, the minimally possible base height measured from FK is 0.05 m when the robot body touches the ground. In the top plot, a position-based RL policy allows the robot to resist the first kick, but the robot falls over after the second kick. In this case, the base height is no longer valid after 1.2 s due to the feet’s loss of contact to the ground. In the bottom plot, with a torque-based RL policy, the robot is able to sustain itself during three successive kicks. Note that the third kick is the hardest one where the quadruped base even contacts the floor. This kicking experiment demonstrates that the torque-based RL policy shows more robustness than the position-based RL policy. We hypothesize that such robustness comes from torque control. Note that this robust behavior is different from the learned recovery behavior in [15], as we do not explicitly train the robot to resist vertical kicks. The robustness of the position-based RL policy can be further improved if we explicitly train with downward kicks.


D. Discussion

1) *Control Frequency*: Our torque-based RL policy is queried at 200 Hz, whereas the query frequency of most position-based RL policies ranges from 20 to 100 Hz [16], followed by a low-level PD controller. During our preliminary test of the policy querying frequency, the GPU on the TX2 can stably infer the 3-layer MLP and maintain the inference at 20 kHz without executing its outputs on the robot hardware, which indicates the potential for a higher-frequency policy if needed. Moreover, compared to the position-based RL policy, the minimum query frequency of the joint torque-based policy is supposed to be higher.


2) *Sim-to-Real Transfer*: One of our concerns in learning torque control was the difficulty of transferring a torque-based policy from simulation to hardware. Surprisingly, it turns out that there is no additional difficulty of sim-to-real transfer for learning torque control, compared to the conventional learning position control. The common domain randomization strategy leads us to successful sim-to-real transfer for both position-based and torque-based policies.

3) *Circumvention of PD Controllers*: Learning torque control allows us to circumvent the use of a low-level PD controller, as the torque-based RL policy directly outputs joint torques. For the position-based locomotion policies, PD gains are usually tuned manually and tightly related to the learning locomotion performance, as studied in [39], [34]. Thus, the torque-based RL method eliminates the number of parameters to tune. Especially, these PD gains need to be tuned on hardware involving extra engineering work.

4) *Potential of Learning Torque Control*: This is the first attempt, to our knowledge, of learning torque control for quadrupedal locomotion with real-world experiments. From a perspective of biomechanics, torque control is more inclined to the way that animals or humans locomote [1], [9]. Hence, we hypothesize that learning torque control can potentially lead to more dynamic and flexible behaviors in the future.



(a) When using the position-based RL policy, the robot is able to sustain itself during the first gentle kick, but after the second larger kick, the robot falls over. The green dashed line stands for the last valid base height measurement before the flip.



(b) When using the torque-based RL policy, the robot is capable of resisting three consecutive kicks successfully, and the third kick is the strongest one during which the robot base even contacts the ground.

Fig. 6: Robustness comparison between learning position control and learning torque control, where the robot base height is illustrated under successive downward kicks while it is standing. The green lines show the measured base heights, and the blue lines depict the nominal base heights. The red lines illustrate the minimal base heights where the base is on the ground. Note that these vertical kicks are not present during training.

IV. CONCLUSION

In this paper, we studied an alternative approach to learning to control a quadruped robot, which uses joint torques as the action space in reinforcement learning, instead of joint positions as conventional learning quadrupedal locomotion work does. We validated the proposed method via real-world experiments, where a quadruped robot is capable of traversing different terrains and resisting external disturbances. Moreover, compared to the learning position control method, we showed that the learning torque control framework can potentially achieve a higher expected return and result in more robust locomotion behaviors. Finally, we provided

discussions on the key features of learning torque control.

ACKNOWLEDGMENTS

We would like to thank Professor Stuart Russell and Professor Pieter Abbeel for their enlightening feedback during the course of this project, and Ayush Agrawal for his help with the experimental setup. We would also like to thank Tingnan Zhang, Xingye Da and Jonathan Li for their insightful discussions on RL-based locomotion. This work was partially supported through a BAIR-Meta AI collaboration and InnoHK of the Government of the Hong Kong Special Administrative Region via the Hong Kong Centre for Logistics Robotics.

REFERENCES

- [1] K. E. Adolph, W. G. Cole, M. Komati, J. S. Garciaguirre, D. Badaly, J. M. Lingeman, G. L. Chan, and R. B. Sotsky, "How do you learn to walk? thousands of steps and dozens of falls per day," *Psychological science*, vol. 23, no. 11, pp. 1387–1394, 2012.
- [2] R. Antonova, S. Cruciani, C. Smith, and D. Kragic, "Reinforcement learning for pivoting task," *arXiv preprint arXiv:1703.00472*, 2017.
- [3] G. Bellegarda and Q. Nguyen, "Robust quadruped jumping via deep reinforcement learning," *arXiv preprint arXiv:2011.07089*, 2020.
- [4] —, "Robust high-speed running for quadruped robots via deep reinforcement learning," *arXiv preprint arXiv:2103.06484*, 2021.
- [5] S. Choi, G. Ji, J. Park, H. Kim, J. Mun, J. H. Lee, and J. Hwangbo, "Learning quadrupedal locomotion on deformable terrain," *Science Robotics*, vol. 8, no. 74, 2023.
- [6] X. Da, Z. Xie, D. Hoeller, B. Boots, A. Anandkumar, Y. Zhu, B. Babich, and A. Garg, "Learning a contact-adaptive controller for robust, efficient legged locomotion," in *Conference on Robot Learning*. PMLR, 2021, pp. 883–894.
- [7] J. Di Carlo, P. M. Wensing, B. Katz, G. Bledt, and S. Kim, "Dynamic locomotion in the mit cheetah 3 through convex model-predictive control," in *2018 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2018, pp. 1–9.
- [8] S. Gangapurwala, M. Geisert, R. Orsolino, M. Fallon, and I. Havoutis, "Rloc: Terrain-aware legged locomotion using reinforcement learning and optimal control," *arXiv preprint arXiv:2012.03094*, 2020.
- [9] D. Han and K. E. Adolph, "The impact of errors in infant development: Falling like a baby," *Developmental science*, vol. 24, no. 5, 2021.
- [10] 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, 2019.
- [11] X. Jiang, H. Wang, Y. Chen, Z. Wu, L. Wang, B. Zou, Y. Yang, Z. Cui, Y. Cai, T. Yu, C. Lv, and Z. Wu, "Mnn: A universal and efficient inference engine," in *MLSys*, 2020.
- [12] Y. Jiang, T. Zhang, D. Ho, Y. Bai, C. K. Liu, S. Levine, and J. Tan, "Simgan: Hybrid simulator identification for domain adaptation via adversarial reinforcement learning," *arXiv:2101.06005*, 2021.
- [13] J. Z. Kolter, P. Abbeel, and A. Y. Ng, "Hierarchical apprenticeship learning with application to quadruped locomotion," in *Advances in Neural Information Processing Systems*. Citeseer, 2008, pp. 769–776.
- [14] A. Kumar, Z. Fu, D. Pathak, and J. Malik, "Rma: Rapid motor adaptation for legged robots," *arXiv preprint arXiv:2107.04034*, 2021.
- [15] J. Lee, J. Hwangbo, and M. Hutter, "Robust recovery controller for a quadrupedal robot using deep reinforcement learning," *arXiv preprint arXiv:1901.07517*, 2019.
- [16] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning quadrupedal locomotion over challenging terrain," *Science robotics*, vol. 5, no. 47, 2020.
- [17] T. Li, N. Lambert, R. Calandra, F. Meier, and A. Rai, "Learning generalizable locomotion skills with hierarchical reinforcement learning," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 413–419.
- [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," *arXiv preprint arXiv:1509.02971*, 2015.
- [19] 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.
- [20] 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, 2022.
- [21] X. B. Peng, M. Andrychowicz, W. Zaremba, and P. Abbeel, "Sim-to-real transfer of robotic control with dynamics randomization," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 3803–3810.
- [22] X. B. Peng, E. Coumans, T. Zhang, T.-W. Lee, J. Tan, and S. Levine, "Learning Agile Robotic Locomotion Skills by Imitating Animals," in *Proceedings of Robotics: Science and Systems*, Corvallis, Oregon, USA, July 2020.
- [23] X. B. Peng and M. 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*, 2017, pp. 1–13.
- [24] L. Pinto, J. Davidson, R. Sukthankar, and A. Gupta, "Robust adversarial reinforcement learning," in *International Conference on Machine Learning*. PMLR, 2017, pp. 2817–2826.
- [25] M. Raibert, K. Blankespoor, G. Nelson, and R. Playter, "Bigdog, the rough-terrain quadruped robot," *IFAC Proceedings Volumes*, vol. 41, no. 2, pp. 10 822–10 825, 2008.
- [26] M. H. Raibert, *Legged robots that balance*. MIT press, 1986.
- [27] D. Reda, T. Tao, and M. van de Panne, "Learning to locomote: Understanding how environment design matters for deep reinforcement learning," in *Motion, Interaction and Games*, 2020, pp. 1–10.
- [28] 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.
- [29] 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.
- [30] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [31] H. Shi, B. Zhou, H. Zeng, F. Wang, Y. Dong, J. Li, K. Wang, H. Tian, and M. Q.-H. Meng, "Reinforcement learning with evolutionary trajectory generator: A general approach for quadrupedal locomotion," *IEEE Robotics and Automation Letters*, 2022.
- [32] J. Siekmann, Y. Godse, A. Fern, and J. Hurst, "Sim-to-real learning of all common bipedal gaits via periodic reward composition," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 7309–7315.
- [33] 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," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 1593–1599.
- [34] L. Smith, I. Kostrikov, and S. Levine, "A walk in the park: Learning to walk in 20 minutes with model-free reinforcement learning," *arXiv preprint arXiv:2208.07860*, 2022.
- [35] Y. Sun, W. L. Ubellacker, W.-L. Ma, X. Zhang, C. Wang, N. V. Csomay-Shanklin, M. Tomizuka, K. Sreenath, and A. D. Ames, "Online learning of unknown dynamics for model-based controllers in legged locomotion," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 8442–8449, 2021.
- [36] 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 *Proceedings of Robotics: Science and Systems*, Pittsburgh, Pennsylvania, June 2018.
- [37] P. Varin, L. Grossman, and S. Kuindersma, "A comparison of action spaces for learning manipulation tasks," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 6015–6021.
- [38] P. Wu, A. Escontrela, D. Hafner, K. Goldberg, and P. Abbeel, "Daydreamer: World models for physical robot learning," *arXiv preprint arXiv:2206.14176*, 2022.
- [39] Z. Xie, X. Da, M. van de Panne, B. Babich, and A. Garg, "Dynamics randomization revisited: A case study for quadrupedal locomotion," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 4955–4961.
- [40] Y. Yang, T. Zhang, E. Coumans, J. Tan, and B. Boots, "Fast and efficient locomotion via learned gait transitions," in *Conference on Robot Learning*. PMLR, 2022, pp. 773–783.
- [41] D. Zhang, A. Loquercio, X. Wu, A. Kumar, J. Malik, and M. W. Mueller, "Learning a single near-hover position controller for vastly different quadcopters," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 1263–1269.