

End-to-End Reinforcement Learning for Torque Based Variable Height Hopping

Raghav Soni^{1,*}, Daniel Harnack^{2,*}, Hauke Isermann^{2,†}, Sotaro Fushimi³, Shivesh Kumar², Frank Kirchner²

Abstract—Legged locomotion is arguably the most suited and versatile mode to deal with natural or unstructured terrains. Intensive research into dynamic walking and running controllers has recently yielded great advances, both in the optimal control and reinforcement learning (RL) literature. Hopping is a challenging dynamic task involving a flight phase and has the potential to increase the traversability of legged robots. Model based control for hopping typically relies on accurate detection of different jump phases, such as lift-off or touch down, and using different controllers for each phase. In this paper, we present an end-to-end RL based torque controller that learns to implicitly detect the relevant jump phases, removing the need to provide manual heuristics for state detection. We also extend a method for simulation to reality transfer of the learned controller to contact rich dynamic tasks, resulting in successful deployment on the robot after training without parameter tuning.

I. INTRODUCTION

Dynamic legged locomotion evolved as a versatile strategy to traverse natural or unstructured terrains. Thus, legged robots such as quadrupeds and humanoids are popular for applications performed in these environments, either autonomously or alongside a human. Quasi-instantaneously making and breaking contacts with the environment is an integral part of legged locomotion, which leads to highly nonlinear, non-smooth dynamics. Thus, from a control perspective, dynamic legged locomotion requires significantly more complex algorithms than e.g. wheeled locomotion. Whereas the problem of dynamic walking on real robots has been solved by various techniques from optimal control (OC) [1], [2], [3] or reinforcement learning (RL) [4], [5], [6], [7], [8], [9], there is considerably less research for the even more dynamic locomotion type of hopping. Hopping can increase a system’s mobility, since it allows for leaping over obstacles that cannot be surpassed otherwise [10], [11]. However, hopping incurs even more control complexity since there is a considerable flight phase during which the system has limited possibilities to adjust for the impact, and the center of mass trajectory is largely determined when the feet leave the ground.

This work was supported by the VeryHuman Project funded by the German Aerospace Center (DLR) with federal funds (FKZ: 01IW20004) and is additionally supported with project funds from the federal state of Bremen for setting up the Underactuated Robotics Lab (Grant Number: 201-342-04-2/2021-4-1)

¹Department of Electronics Engineering, Indian Institute of Technology (Banaras Hindu University), Varanasi, India.

²DFKI GmbH Robotics Innovation Center, Bremen, Germany

³Undergraduate Course Program of Mechanical and Systems Engineering, Kyoto University, Kyoto, Japan

[†]Corresponding author: Hauke.Isermann@dfki.de

*Both authors contributed equally




Fig. 1: RL based torque controlled jumping snapshots in simulation and on the real robot.

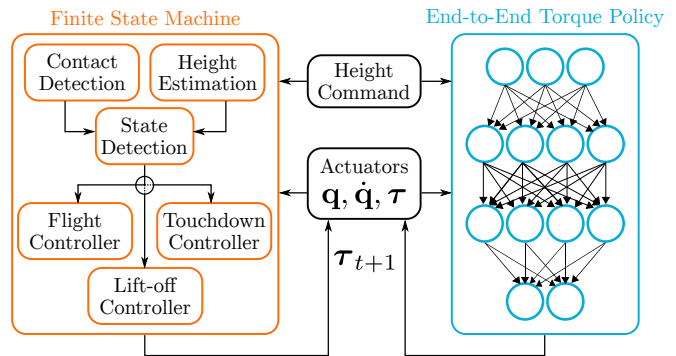


Fig. 2: Comparison of control concepts.

The canonical system to study hopping, which is also used in this paper, is a single hopping leg or monopod. Indeed, one of the earliest robotic systems showing dynamic legged locomotion was a single leg that could navigate a flat surface by jumping [12]. To control the hopping height, a heuristically tuned force controller was used, motivated by an energy shaping algorithm [13]. These seminal studies sparked a wealth of research into the control of single legged hopping machines. A common theme of such controllers is the reliance on detection of various states, e.g. lift-off, peak attitude, touchdown, and minimum attitude [13], [14], [15], [16], [17], [18]. The full jumping controller is then realized as a state machine, where PD controllers are typically used during flight phases, while stance phase states are directly controlled by torque or force. Deploying these controllers on hardware requires hand tuning parameters and system specific adaptations to account for model inaccuracies or unmodelled dynamics. Also, the detection of different jump states and appropriate control output during the lift-off phase relies on accurate height estimation and contact detection, for which further heuristics are typically employed.

RL offers the promise to alleviate these issues. We hypothesize that, since neural networks are universal function approximators [19], [20], learning based controllers with neural

network function approximators are able to implicitly detect relevant jump phases and thus realize a unified hopping controller without the need for explicit state transition heuristics. Even more, a fully integrated end-to-end solution for hopping should be possible via RL, mapping only proprioceptive feedback, i.e. actuator positions and velocities, to direct torque control, since this proprioceptive data is theoretically sufficient to implement a variable jumping height controller. The comparison of our method to a classical approach is visualized in Fig. 2.

Whereas several previous studies utilized learning based controllers for jumping [11], [21], they focussed on single leaps, making implicit detection of jump phases less critical than for continuous jumping. From the data, it can also not be derived whether such a phase detection is actually realized. Recently, a RL based continuous hopping controller with adjustable jumping height for a small quadruped was developed in [22]. While this shows the feasibility of a unified controller that does not require a state machine, it still relied on height estimation and PD control, and thus can not be considered as a truly end-to-end learning approach. The widespread use of PD controllers in RL research is in part a consequence of the low sample complexity of many algorithms, which necessitates training in simulation. Using PD controllers reduces the requirements on the accuracy of the dynamics simulation and thus increases the chances of successful simulation to reality transfer. However, this requires tuning of PD gains for a successful sim2real transfer and may additionally hinder performance in highly dynamic tasks such as jumping, where direct torque control can unlock the full dynamical capabilities of a system [4], [23].

In summary, all previous approaches from classical control and RL require a subset of height estimation, contact detection, hyperparameter tuning, PD control, or a behavior state machine. In this paper, we present a RL based method that requires none of the above. We show successful training and simulation to reality transfer of a torque controller with implicit jumping phase detection and controllable jumping height, while only relying on proprioceptive feedback¹. To achieve this, we draw inspiration from energy shaping for the design of the reward function, and extend a previous technique for accurate simulation to reality transfer [24] to higher dimensional parameter spaces and dynamic, contact rich tasks. To the best knowledge of the authors, such a controller is described for the first time for a monopod.

II. MATERIALS AND METHODS

A. Robotic System

The robot used for the experiment is a custom-made 3 degrees of freedom (DOF) hopping leg system, mounted on a vertical rail with 1 passive DOF and 2 active DOFs. Fig. 3 shows a photo of the system, along with a 3D design model. The 2 active DOFs in the leg are actuated via quasi-direct drive motors qdd100 from mjbots [25] operating at a




Fig. 3: Hopping leg used in the experiments.

frequency of $200Hz$. While the shoulder joint shares the joint axis with the motor axis, the elbow joint is driven by a motor placed at the shoulder via a belt drive with transmission ratio = 1:2. The housing is a light weight carbon fiber construction.

B. Energy Shaping

As a baseline comparison, we implemented a classical energy shaping (ES) controller in simulation [26]. This ES controller is part of a finite state machine [27]. As shown in Fig. 4 the state machine consists of three states:

- *Lift-off*: This phase is used to apply the desired energy with the ES controller for the next jump. It ends when the leg loses its ground contact.
- *Flight*: In the flight phase, the leg prepares for the touchdown by moving into a predefined pose. The phase ends with the first contact of the leg with the ground.
- *Touchdown*: During touchdown, the leg damps its movement using high damping and low positional gains. It ends when the base velocity $\dot{x} \geq 0$.

1) *Controller design*: As mentioned above, the desired energy E_d which is required to jump to a desired base height x_d has to be applied during the lift-off phase. For simplicity, we assume the robot to be a point mass m at its base. In this case, the required energy can be calculated with:

$$E_d = mgx_d \quad (1)$$

Here, g is the gravitational acceleration.

Accordingly, we can estimate the reached energy E_{j-1} from the last jump using the estimated jumping height x_{j-1} :

$$E_{j-1} = mgx_{j-1} \quad (2)$$

For reaching this energy, we needed to apply a feed-forward force $F_{f,j-1}$ to the ground while the leg had ground contact. For the next jump, we estimate the new feed-forward force with:

$$F_{f,j} = \frac{mg(x_d - x_{0,j})}{\Delta x_{l,j}} \quad (3)$$




Fig. 4: State machine used for the energy shaping controller.

¹https://github.com/dfki-ric-underactuated-lab/hopping_leg

Here, $x_{0,j}$ is the current minimum base height after the touchdown, and $\Delta x_{l,j}$ is the the expected distance to be covered during the current lift-off, calculated from $x_{0,j}$. Since the formerly applied force is proportional to the reached energy E_{j-1} we can write:

$$E \propto kF_f \quad (4)$$

and as the feed-forward force is almost constant:

$$\dot{E} \propto \dot{k}F_f \quad (5)$$

Thus, we can control the energy in the system by altering the gain k . To control the gain k the following update rule has been used:

$$k_j = k_{j-1} \left(\frac{E_d}{E_{j-1}} \right)^2 \quad (6)$$

As joint torque controller, Cartesian stiffness control with the following control law was used [28]:

$$\tau = J^T(q) \left[k_j \begin{pmatrix} F_{f,j} \\ 0 \end{pmatrix} + k_{p,y} \begin{pmatrix} 0 \\ -y \end{pmatrix} + k_{d,y} \begin{pmatrix} 0 \\ -\dot{y} \end{pmatrix} \right] \quad (7)$$

Here, τ is the vector of desired joint torques, $J^T(q)$ is the transpose of the hybrid Jacobian at the end-effector and $K_{p,y}$ and $K_{d,y}$ are the Cartesian PD gains in y direction (refer to Fig. 3). As shown in (7), the PD terms of the Cartesian stiffness controller are only responsible to maintain the y position of the end-effector, while the energy shaping control is used to apply the forces in x direction.

2) *Jumping height estimation:* For the ES controller, a height feedback is necessary. Therefore, a proprioceptive height estimation has been implemented. During flight phase, no additional forces can be applied to the system. Hence, we can expect the base acceleration \ddot{x} to be roughly equal to the gravitational acceleration g . Due to the proprioceptive feedback, we know the lift-off position x_l and velocity \dot{x}_l . Thus, the current base height x during flight phase can be calculated with:

$$x = \frac{1}{2}g(t^2 - t_l^2) - g t_l(t - t_l) + \dot{x}_l(t - t_l) + x_l \quad (8)$$

Here, t is the current time and t_l is the time at lift-off. During stance phase, the current height is calculated from forward kinematics.

3) *Simulation and real system parameters:* The simulations of the ES controller have been performed using PyBullet physics simulation [29]. For the simulation a control frequency of 400Hz has been used to control the joint torques. The parameters $k_{j=0} = 1.0$, $k_{p,y} = 10.0$, $k_{d,y} = 3.0$ were optimized manually. For the real system, the control frequency has been reduced to 200Hz .

C. Reinforcement Learning

1) *Problem Formulation:* The hopping leg problem is formulated as a Markov Decision Process (MDP), where the agent, i.e. the controller in this case, interacts with the environment, i.e. the leg and its surroundings. A MDP is given by a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R})$, where \mathcal{S} is the set of states called the state space, \mathcal{A} is the set of actions called the action

space, $\mathcal{P}(s_{t+1} | s_t, a_t)$ the probability that taking action a_t in state s_t will lead to state s_{t+1} , and $\mathcal{R}(s_t, a_t, s_{t+1})$ the expected immediate reward for transitioning from s_t to s_{t+1} by taking the action a_t . At each time step t , an action $a_t \sim \pi(a_t | s_t)$ is sampled from the policy given the current state s_t . The objective of RL is to optimize the policy π such that the expected return is maximized.

From the variety of RL algorithms, we choose Soft Actor-Critic (SAC) [30], a state-of-the-art off-policy algorithm, since it is relatively sample-efficient, stable, and requires little to no hyperparameter tuning. SAC aims to maximize the expected reward while also maximizing the policy entropy H . The objective is formulated as

$$\pi^* = \arg \max_{\pi} \mathbb{E}_{a \sim \pi} \left[\sum_{t=0}^{\infty} \gamma^t \left(R(s_t, a_t, s_{t+1}) + \alpha H(\pi(\cdot | s_t)) \right) \right].$$

Maximizing the entropy as a secondary objective leads to policies that are maximally variable while performing the task, making them intrinsically robust.

For continuous actions, exploration is commonly done in action space. At each time step, a noise vector ϵ_t is sampled from a Gaussian distribution and added to the action output, such that $\pi(a_t | s_t) \sim \mu(s_t, \theta_\mu) + \mathcal{N}(0, \sigma^2)$, where μ is the deterministic policy and θ_μ its parameters. We use the modification of generalized state dependent exploration (gSDE) [31]. Here, the noise vector is a function of the state and the policy features $z_\mu(s_t, \theta_{z_\mu})$, which is the last layer before the deterministic output $\mu(s_t) = \theta_\mu z_\mu(s_t, \theta_{z_\mu})$, i.e. $\epsilon_t(s_t, \theta_\epsilon) = \theta_\epsilon z_\mu(s_t)$. With gSDE, the action for a given state s_t remains the same until the noise parameters are sampled again. This promotes more consistent exploration and results in reduced shaky behavior on hardware [31].

2) *Network architecture:* The policy is modeled with a multilayer perceptron (MLP) with four hidden layers of 256, 256, 128, and 128 neurons. The activation function is ReLU. The critic network is modeled by a separate network with the same architecture. LSTM policy networks were also tried but offered no empirical advantage. The policy is inferred at the operating frequency of 200Hz .

3) *Observation and action space:* The hopping leg system has no additional sensors apart from the joint encoders. Thus, only normalized joint positions and velocities, and the desired jumping height over the three last time-steps t , $t-1$, and $t-2$ constitute the observation state. Joint data over multiple time-steps is empirically found to be essential to produce the desired behaviour with implicit contact detection. Hence, the observation space is $s \in \mathbb{R}^{3 \times 5 = 15}$. The action space consists of the normalized output motor torques, which are later scaled up before being sent as the torque commands. The action space is thus $a \in \mathbb{R}^2$.

4) *Reward:* The total reward at each time step is a weighted sum of positive gains and negative penalties, encoding behaviours to be encouraged or precluded. The reward comprises the following components:

a) *Energy Gain (G_e):* The agent is incentivized to maximize the kinetic and elastic potential energy of the leg at any given time step. The reasoning behind this term is

an approximation of the leg by spring with mean length x_o . This reward term promotes an oscillatory behaviour leading to high enough velocities for hopping. The corresponding term is calculated as:

$$G_e = \dot{x}^2 + (x - x_o)^2 \quad (9)$$

where x and \dot{x} are the base height and velocity, respectively. x_o is the base height for the initial standing position of the leg.

b) *Height barrier penalty (P_h)*: The agent is penalized exponentially when the base height crosses the desired height command x^d .

$$P_h = \begin{cases} 1 - e^{x-x^d}, & \text{if } x \geq x^d \\ 0, & \text{otherwise} \end{cases} \quad (10)$$

c) *Jerky Action Penalty (P_j)*: Sudden changes in the output torques can cause shakiness in the hardware, making the policy hard to transfer. Therefore, the agent is penalized for large differences in consecutive actions.

$$P_j = \sum_{i=0}^2 (a_t^i - a_{t-1}^i)^2 \quad (11)$$

d) *Joint constraints penalty (P_{jp}, P_{jv})*: It is desired to keep the joint position limits within some pre-defined constraints to avoid self-collisions and prevent arbitrary configurations. The joint velocities should be reasonably bounded for successful sim-to-real transfer. These constraints are imposed with negative penalties. The penalty for the position limit is structured such that it becomes significant around the limits and beyond them but stays reasonably low elsewhere. It is calculated as:

$$P_{jp} = \sum_{i=0}^2 \begin{cases} e^{-10(q_i - q_i^l)} + e^{10(q_i - q_i^h)}, & \text{if } q_i^l \leq q_i \leq q_i^h \\ 1, & \text{otherwise} \end{cases} \quad (12)$$

Here, q_i^l and q_i^h denote the lower and upper joint limits, respectively. To reasonably constrain the search space for the agent, we used a PD controller to bring the joints back within bounds if joint limits are passed during training. The joint velocities are penalized if they cross the saturation limits for the motors.

$$P_{jv} = \sum_{i=0}^2 \begin{cases} 0, & \text{if } -\dot{q}_i^h \leq \dot{q}_i \leq \dot{q}_i^h \\ \dot{q}_i^2 - \dot{q}_i^h, & \text{otherwise} \end{cases} \quad (13)$$

Here, \dot{q}_i^h is the maximum desired joint velocity.

The final expected reward is calculated as:

$$R = w_1 G_e - w_2 P_h - w_3 P_j - w_4 P_{jp} - w_5 P_{jv} \quad (14)$$

The weights used during training are $w_1 = 0.5$, $w_2 = 2$, $w_3 = 0.05$, $w_4 = 0.02$, and $w_5 = 0.005$.

D. Simulation to Reality Transfer

We use a custom gym [32] environment with MuJoCo physics engine [33] for training in simulation. As explored in [34], MuJoCo is well suited for robotics and reinforcement learning problems as it provides a wide range of solver




Fig. 5: Parameter optimization pipeline.

parameters and settings, which can be adapted and optimized for many use cases. The policies trained in MuJoCo with default model and simulation parameters failed to transfer to the hardware. Therefore, we optimised the simulation parameters to narrow the sim-to-real gap.

1) *Simulation Parameter Optimisation*: The goal of this step is to match simulation dynamics to the real robot using simple training trajectories.

a) *Trajectory Generation and Data Collection*: A varied set of task-space, hand-tuned sinusoidal trajectories is generated for two different system configurations. These include a fixed-base configuration, where the leg is suspended in the air, and a moving-base configuration, where the leg comes in contact with the ground. For the fixed-base configuration, the template trajectories are:

$$\begin{aligned} x &= (A \cos\left(\frac{2\pi t}{T_1}\right) + \epsilon) \cos(\theta) \\ y &= (A \cos\left(\frac{2\pi t}{T_1}\right) + \epsilon) \sin(\theta) \\ \theta &= -\frac{\pi}{2} \cos\left(\frac{2\pi t}{T_2}\right) \end{aligned} \quad (15)$$

Here, $\epsilon = L_1 + L_2 - A$ is the trajectory offset from the origin, given L_1 and L_2 are the shank and calf link lengths for the leg. The periods (T_1, T_2) and the amplitude A are varied such that the maximum workspace of the leg is covered. The moving-base trajectory only consists of vertical trajectories with a few segments fast enough to break ground contact. These trajectories closely imitate the hopping configuration of the leg and are given as:

$$\begin{aligned} x &= A \cos\left(\frac{2\pi t}{T_1}\right) + \epsilon \\ y &= 0 \end{aligned} \quad (16)$$

The joint-level trajectories are obtained through inverse kinematics and tracked on the hardware with a PD controller

running at a frequency of 200Hz . The controller gains are fixed and the target velocity set to 0. For both fixed-base and floating base configuration, 240 s of data was recorded where $T_1 \in \{0.75, 0.5, 0.25\}$, $T_2 \in \{10, 20\}$, and $A \in \{0.15, 0.1, 0.05\}$. Joint positions, velocities, and resulting motor torques are recorded on the actual hardware.

b) Simulation Parameters Optimisation: Using the hardware trajectories, we optimize for the simulation’s dynamics and solver parameters. We use the same PD controller running at the same frequency to track the generated trajectories in simulation, with gains adjusted for the motors’ internal gear ratio. We optimize for the following set of simulation parameters:

- *Dynamic Parameters:* The simulation model’s dynamic parameters to be optimized involve the friction loss, damping and armature (rotor inertia) values for the hip and knee motors, friction loss and damping for the rail, which is modeled by a passive prismatic joint, and the link inertias.
- *Solver Parameters:* Time constant and damping ratio are two of the solver parameters, characteristic of the mass-spring-damper constraint modeling of MuJoCo. These parameters are optimized to modulate the contact model between the leg and the plane.

CMA-ES [35] is used to optimize these parameters with a cost on the cumulative joint position difference between simulation trajectories and recorded real hardware data at each time step.

$$J(q) = \sum_{t=0}^{t_f} \sum_{i=0}^1 (q_{i_{\text{sim}}}^t - q_{i_{\text{real}}}^t)^2 \quad (17)$$

In high dimensional optimization problems, such as this, it can become hard for the solver to converge and find an optimal solution. Dynamic coupling also occurs between the parameters, potentially leading to low cost but poor transfer to the hardware. To prevent these issues, the parameters that can be roughly estimated during modeling, i.e. link inertia and solver parameters, are fixed in a first optimization pass. In a second pass, we optimize all dynamic and solver parameters while placing bounds on the friction, damping, and armature parameters derived from the first pass. This helps converging to an optimal and pragmatic solution. Fig. 5 visualizes the optimization procedure.

2) Policy Robustness: Two methods are employed during training to make the policy robust to delays, noise, and other disturbances.

a) Delays: As mentioned before, the observation space consists of sensor readings from the last three consecutive time steps. In addition, the observation data over the last ten time-steps is stored in a buffer. While training, with a probability of 0.5 at each time step, data of three time steps is randomly sampled from the buffer in correct temporal sequence and used as observation instead. This helps to simulate plausible delays on the real system, effectively making the policy more robust.

TABLE I: The simulation parameters obtained after CMA-ES optimization.

Joint	Friction loss	Damping	Armature
Rail Prismatic Joint	0.7024	1.0724	-
Hip Joint	0.4364	0.0005	0.00004
Knee Joint	0.0015	0.1441	0.0001

Parameter	Value
Hip Link Z Inertia	0.004061
Knee Link Z Inertia	0.000845
Time Constant	0.0911
Damping Ratio	0.6678

b) Noise: Noise is added to the joint data and the torques given by the policy to simulate sensor noise and control inaccuracies. At each time step, the noise is sampled from a uniform distribution ranging from $-\lambda u$ to λu , where λ is the error range and u is the observed value. We set $\lambda = 0.05$ for the joint positions and velocities, and $\lambda = 0.15$ for the output torques.

III. RESULTS

The optimization of simulation parameters described in Section II lead to the parameters shown in Table I. Training for the jump heights 0.25 m, 0.3 m, and 0.35 m in simulation yielded a controller that is able to interpolate between these three desired jump heights, showing that the controller learned an approximation of the task space inverse dynamics for this problem. Figure 6 shows the jump height of a 30 s trial with an initial desired jump height of 0.25 m. Every 5 s the desired jump height is increased by 0.02 m. While there is a significant deviation of the actual average jump height especially for intermediate commands, the mapping from desired to actual jump heights is monotonic.

To assess the implicit contact detection of the controller, we analyse the applied torque in simulation to the elbow joint for a 10 s trial with a commanded jump height of 0.30 m. Figure 7 shows the controller torque output for all encountered configurations in the phase space of the actuated elbow joint. It is evident that during the stance phase, the controller applies significantly higher torques than in the flight phase, to generate the lift-off. In addition, the control torque increases after the minimal attitude is reached. This strongly implies that the controller indeed detects ground contact, solely based on the proprioceptive observation of joint positions and velocities.

The controller trained in simulation was tested on the real robot without further adjustment. Figure 8 shows the base height trajectories for simulated and real robot for a trial with changing desired jump height. Whereas the real jump height is lower than the commanded height, the ordering of jump heights is as intended, i.e. a higher desired height leads to a higher actual jump height. The offset between commanded and actual jump height lies between 0.04 m and 0.06 m.




Fig. 6: Jumping heights of the simulated robot in a 30 second trial with monotonically increasing desired jump heights. The desired heights increase from 0.25 m to 0.35 m in increments of 0.02 m. The controller was trained with only three desired jump heights of 0.25 m, 0.30 m, and 0.35 m.




Fig. 7: Controller torque output during 10 s of continuous jumping with a commanded height of 0.30 m in phase space of the elbow joint. The black (grey) line is an example phase space trajectory for one flight (stance) phase. Circles and diamonds denote starting and end points of the jump phase.

We observe that the temporal structure of consecutive jumps differs between simulation and the real robot. For a desired jump height of 0.25 m, the real robot shows jump heights alternating between ≈ 0.2 and ≈ 0.24 m. For a commanded height of 0.35 m, the jump frequency on the real robot is slightly reduced, because two jumps, around 11 and 12.5 s, were 'skipped'. Note that the absolute values of the jump heights on the real system are not exact, since it is determined by tracking the center of the upper motor in video recordings. This induces some noise on the measurement. In addition, a height dependent small parallax error can be expected.

For a statistical analysis of the jump height distribution for varying height commands, the data of the 15 s trial shown




Fig. 8: Jumping heights of simulated and real robot for 15 seconds. The commanded jump height is 0.25 m for the first 5 s, 0.30 m for the next 5 s, 0.35 m for the last five seconds. Both in simulation and the real robot, increasing the desired jump height leads to higher actual jumps.




Fig. 9: Jump height distributions for different commands, in simulation and on the real system. Colored lines show the medians of the distributions, gray lines the respective commanded height. In simulation, the distributions for the RL controller are completely non-overlapping, thus clearly significantly different. On the real system, the distributions show some overlap. For comparison, the distributions of jump heights generated by an ES controller are also shown.

in Figure 8 is merged with 15 s trials of fixed jump heights at 0.25 m, 0.30 m, and 0.35 m, both in simulation and the real robot (also see the accompanying video available as the multimedia attachment). Figure 9 show the resulting jump height distributions for the trained controller, along with simulation results for an energy shaping controller as reference. While the distributions of RL based jumping are completely separated for the simulation, the real system tests show overlaps between neighbouring jump height distribu-

tions. This can at least partly be attributed to the noise induced by the pixel tracking used to estimate the heights in the real experiments. The bimodal nature of the distribution for a commanded height of 0.25 m is a consequence of the alternating jump heights, also seen in Figure 8. We use a Wilcoxon-Ranksum test to evaluate the difference between neighbouring distributions. Both in simulation and on the real system, all neighbouring distributions are significantly different at $p < 0.001$.

Our baseline energy shaping controller worked remarkably well in both simulation and on real hardware (see Fig. 9). This is as expected since we exploit the model knowledge and physics that captures the essence of jumping task. However, it requires expert knowledge to tune the contact detection threshold and other controller gains which can be time consuming. Our proposed End-to-End RL controller does not require such expert knowledge and demonstrates a similar trend for different jumping heights. The standard deviation in jumping height is even smaller in some cases, especially in simulation. However, there is substantial room for improvement in the performance of the RL controller on the real system in comparison to the baseline ES controller.

IV. DISCUSSION

The main objective was to find a jumping controller mapping proprioceptive feedback to torque control, including the avoidance of height estimation and PD control strategies as used by [22]. Thus it may seem counterintuitive that we impose soft joint limits with a PD controller and use the base height for reward calculation. However, since variable height jumping cannot be defined without the notion of height, it is strictly necessary information for the agent such that the task space inverse dynamics can be approximated. We want to emphasize though that the height is only used in the reward during training, and is not required as direct feedback to the controller. The soft joint limits serve as a gentle exploration guiding strategy, similar to initial example trajectories as used by [22]. On the hardware, they are still in place for safety reasons, but are rarely crossed. Thus, our control approach can be considered truly end-to-end. In the following, further features and critical design decisions are discussed in more depth.

A prerequisite for successful transfer to the real system is a small simulation to reality gap. Prior to developing the current approach, the more common technique of domain randomization [36] was also tested, which generated unsatisfying behavior transfer. Our approach is adapted from [24], who used simulation parameter optimization to make trajectories in simulation follow real recorded training data. We extend this approach by introducing a two stage process for high dimensional parameter spaces and showing the applicability to collision rich and dynamic tasks. We chose to keep inertia parameters fixed in the first stage, since they can be reasonably well estimated from the structure, whereas other dynamical parameters are much harder to infer a priori. We also noted that identifying the rotor inertia was crucial. While not necessary for less dynamic behavior such

as walking, rotor inertia becomes more influential for highly dynamic motions. The superior simulation to reality transfer can be explained by domain randomization leading to a trade off between generality over a range of parametrizations to optimality on the actual hardware, which has well defined parameters. In contrast, the method we propose is more akin to dynamic system identification. However, the target is not the true physical parameters, but the closest possible representation of the system dynamics within the simulation.

To make the policy robust to expectable delays on the real system, we used random sampling from an observation buffer. This random sampling technique is easy to implement without having to know the exact delays and their distributions. An alternative approach would be to use an actuator model as suggested by [5] to learn quadruped walking. In their case, a good motor model was probably more relevant since the robot’s legs use series elastic actuators, which are expected to have more complex delay dynamics. If this is not the case, we argue for our method as a simpler solution.

The policy shows good interpolation performance for height values that were not explicitly included in the training. This suggests that the policy implicitly learned a task space inverse dynamics model of the system. This assumption is further supported by the implicit detection of different jump phases. However, height tracking shows relatively higher deviations at intermediate commands around 0.3 m. This could be an issue of the neural network not having enough capacity to represent the full dynamics. A thorough hyperparameter tuning of the network architecture could improve the results, but is out of scope for this paper.

The remaining differences in the jump heights between simulation and reality can be a consequence of non-optimal dynamic parameters of the simulator. However, we noted that adding more data to the parameter optimization pipeline did not significantly change the optimization result. Another explanation could be additional, unmodelled non-linear dynamics such as motor backlash, motor torque saturation, or state dependent sensor noise. A strategy to improve performance without having to explicitly model these effects is to continue training the controller on the real system directly, using the current policy as a starting point. For this, the used SAC algorithm is particularly well suited [37].

V. CONCLUSION

In summary, we presented a method to train a unified torque controller for continuous hopping with a monoped robot. The controller is able to interpolate between jump heights and implicitly detect relevant jump phases and act accordingly. The simulation to reality mapping procedure eliminates the need of parameter tuning for behavior transfer. The trained policy realizes a direct mapping from proprioceptive feedback to torque control. To the authors’ knowledge, this is the first reported end-to-end training procedure for a jump height adjustable monoped torque controller. However, much needs to be done to bring the height tracking accuracy of this approach closer to the model based energy shaping control. Future research directions include a thorough hyperparameter

tuning of the neural network architecture to improve the jump height interpolation in simulation, as well as continued training on the real system to mitigate the effect of residual dynamics modeling inaccuracies of the simulator. We also plan to integrate this work in the RealAIGym ecosystem [38] similar to other canonical underactuated systems like simple pendulum [39], double pendulum [40], and AcroMonk [41].

REFERENCES

- [1] K. Mombaur, J.-P. Laumond, and E. Yoshida, "An optimal control model unifying holonomic and nonholonomic walking," in *Humanoids 2008-8th IEEE-RAS International Conference on Humanoid Robots*. IEEE, 2008, pp. 646–653.
- [2] H. Osumi, S. Kamiya, H. Kato, K. Umeda, R. Ueda, and T. Arai, "Time optimal control for quadruped walking robots," in *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006*. IEEE, 2006, pp. 1102–1108.
- [3] A. Herd, E. A. Cousineau, C. M. Hubicki, and A. D. Ames, "3d dynamic walking with underactuated humanoid robots: A direct collocation framework for optimizing hybrid zero dynamics," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 1447–1454.
- [4] S. Chen, B. Zhang, M. W. Mueller, A. Rai, and K. Sreenath, "Learning torque control for quadrupedal locomotion," *arXiv preprint arXiv:2203.05194*, 2022.
- [5] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter, "Learning agile and dynamic motor skills for legged robots," *Science Robotics*, vol. 4, no. 26, p. eaau5872, 2019.
- [6] Z. Li, X. Cheng, X. B. Peng, P. Abbeel, S. Levine, G. Berseth, and K. Sreenath, "Reinforcement learning for robust parameterized locomotion control of bipedal robots," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 2811–2817.
- [7] D. Jain, A. Iscen, and K. Caluwaerts, "Hierarchical reinforcement learning for quadruped locomotion," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 7551–7557.
- [8] Z. Xie, P. Clary, J. Dao, P. Morais, J. Hurst, and M. Panne, "Learning locomotion skills for cassie: Iterative design and sim-to-real," in *Conference on Robot Learning*. PMLR, 2020, pp. 317–329.
- [9] G. B. Margolis, G. Yang, K. Paigwar, T. Chen, and P. Agrawal, "Rapid locomotion via reinforcement learning," *arXiv preprint arXiv:2205.02824*, 2022.
- [10] E. W. Hawkes, C. Xiao, R.-A. Peloquin, C. Keeley, M. R. Begley, M. T. Pope, and G. Niemeyer, "Engineered jumpers overcome biological limits via work multiplication," *Nature*, vol. 604, no. 7907, pp. 657–661, 2022.
- [11] G. Bellegarda and Q. Nguyen, "Robust quadruped jumping via deep reinforcement learning," *arXiv preprint arXiv:2011.07089*, 2020.
- [12] M. H. Raibert, H. B. Brown Jr, and M. Chepponis, "Experiments in balance with a 3d one-legged hopping machine," *The International Journal of Robotics Research*, vol. 3, no. 2, pp. 75–92, 1984.
- [13] M. H. Raibert, "Hopping in legged systems—modeling and simulation for the two-dimensional one-legged case," *IEEE Transactions on Systems, Man, and Cybernetics*, no. 3, pp. 451–463, 1984.
- [14] G. J. Zeglin, "Uniroo—a one legged dynamic hopping robot," Ph.D. dissertation, Massachusetts Institute of Technology, 1991.
- [15] K. Harbick and G. S. Sukhatme, "Controlling hopping height of a pneumatic monopod," in *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No. 02CH37292)*, vol. 4. IEEE, 2002, pp. 3998–4003.
- [16] A. Lebaudy, J. Presser, and M. Kam, "Control algorithms for a vertically-constrained one-legged hopping machine," in *Proceedings of 32nd IEEE Conference on Decision and Control*. IEEE, 1993, pp. 2688–2693.
- [17] G. Zhao, F. Szymanski, and A. Seyfarth, "Bio-inspired neuromuscular reflex based hopping controller for a segmented robotic leg," *Bioinspiration & biomimetics*, vol. 15, no. 2, p. 026007, 2020.
- [18] J. Ramos, Y. Ding, Y.-w. Sim, K. Murphy, and D. Block, "Hoppy: an open-source and low-cost kit for dynamic robotics education," *arXiv preprint arXiv:2010.14580*, 2020.
- [19] A. Pinkus, "Approximation theory of the mlp model in neural networks," *Acta numerica*, vol. 8, pp. 143–195, 1999.
- [20] Z. Lu, H. Pu, F. Wang, Z. Hu, and L. Wang, "The expressive power of neural networks: A view from the width," *Advances in neural information processing systems*, vol. 30, 2017.
- [21] Y. Kuang, S. Wang, B. Sun, J. Hao, and H. Cheng, "Learning jumping skills from human with a fast reinforcement learning framework," in *2018 IEEE 8th Annual International Conference on CYBER Technology in Automation, Control, and Intelligent Systems (CYBER)*. IEEE, 2018, pp. 510–515.
- [22] M. Bogdanovic, M. Khadiv, and L. Righetti, "Model-free reinforcement learning for robust locomotion using trajectory optimization for exploration," *arXiv preprint arXiv:2107.06629*, 2021.
- [23] D. M. Wolpert, J. Diedrichsen, and J. R. Flanagan, "Principles of sensorimotor learning," *Nature reviews neuroscience*, vol. 12, no. 12, pp. 739–751, 2011.
- [24] M. Kaspar, J. D. M. Osorio, and J. Bock, "Sim2real transfer for reinforcement learning without dynamics randomization," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 4383–4388.
- [25] mjbots, "qdd100 beta 3 servo," <https://mjbots.com/collections/servos-and-controllers/products/qdd100-beta-3>, accessed: 10-09-2022.
- [26] R. Tedrake, *Underactuated Robotics*, 2023. [Online]. Available: <https://underactuated.csail.mit.edu>
- [27] S. Hyon, T. Emura, and T. Mita, "Dynamics-based control of a one-legged hopping robot," *Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering*, vol. 217, no. 2, pp. 83–98, 2003.
- [28] J. Di Carlo, "Software and control design for the mit cheetah quadruped robots," Ph.D. dissertation, Massachusetts Institute of Technology, 2020.
- [29] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," <http://pybullet.org>, 2016.
- [30] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," 2018. [Online]. Available: <https://arxiv.org/abs/1801.01290>
- [31] A. Raffin and F. Stulp, "Generalized state-dependent exploration for deep reinforcement learning in robotics," *Arxiv*, 2020.
- [32] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "Openai gym," 2016.
- [33] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012, pp. 5026–5033.
- [34] M. Körber, J. Lange, S. Rediske, S. Steinmann, and R. Glück, "Comparing popular simulation environments in the scope of robotics and reinforcement learning," 2021. [Online]. Available: <https://arxiv.org/abs/2103.04616>
- [35] N. Hansen and A. Ostermeier, "Completely derandomized self-adaptation in evolution strategies," *Evolutionary Computation*, vol. 9, no. 2, pp. 159–195, 2001.
- [36] 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.
- [37] T. Haarnoja, S. Ha, A. Zhou, J. Tan, G. Tucker, and S. Levine, "Learning to walk via deep reinforcement learning," *arXiv preprint arXiv:1812.11103*, 2018.
- [38] F. Wiebe, S. Vyas, L. J. Maywald, S. Kumar, and F. Kirchner, "RealAIGym: Education and Research Platform for Studying Athletic Intelligence," in *Proceedings of Robotics Science and Systems Workshop Mind the Gap: Opportunities and Challenges in the Transition Between Research and Industry*, New York, July 2022.
- [39] F. Wiebe, J. Babel, S. Kumar, S. Vyas, D. Harnack, M. Boukheddimi, M. Popescu, and F. Kirchner, "Torque-limited simple pendulum: A toolkit for getting familiar with control algorithms in underactuated robotics," *Journal of Open Source Software*, vol. 7, no. 74, p. 3884, 2022. [Online]. Available: <https://doi.org/10.21105/joss.03884>
- [40] F. Wiebe, S. Kumar, L. Shala, S. Vyas, M. Javadi, and F. Kirchner, "An open source dual purpose acrobat and pendubot platform for benchmarking control algorithms for underactuated robotics," *IEEE Robotics and Automation Magazine*, 2023, under review.
- [41] M. Javadi, D. Harnack, P. Stocco, S. Kumar, S. Vyas, D. Pizzutilo, and F. Kirchner, "Acromonk: A minimalist underactuated brachiating robot," *IEEE Robotics and Automation Letters*, vol. 8, no. 6, pp. 3637–3644, 2023.