

RESEARCH ARTICLE

Learning Bipedal Walking for Humanoids With Current Feedback

ROHAN P. SINGH^{1,2}, ZHAOMING XIE³, PIERRE GERGONDET¹,
AND FUMIO KANEHIRO^{1,2}, (Member, IEEE)

¹CNRS-AIST JRL (Joint Robotics Laboratory), IRL, National Institute of Advanced Industrial Science and Technology (AIST), Tsukuba, Ibaraki 305-8560 Japan

²Department of Intelligent and Mechanical Interaction Systems, Graduate School of Science and Technology, University of Tsukuba, Ibaraki 305-8577, Japan

³Department of Computer Science, Stanford University, Stanford, CA 94305, USA

Corresponding author: Rohan P. Singh (rohan-singh@aist.go.jp)

This work was supported in part by the Japan Science and Technology (JST) SPRING Fellowship Program under Grant JPMJSP2124.

ABSTRACT Recent advances in deep reinforcement learning (RL) based techniques combined with training in simulation have offered a new approach to developing robust controllers for legged robots. However, the application of such approaches to real hardware has largely been limited to quadrupedal robots with direct-drive actuators and light-weight bipedal robots with low gear-ratio transmission systems. Application to real, life-sized humanoid robots has been less common arguably due to a large *sim2real* gap. In this paper, we present an approach for effectively overcoming the *sim2real* gap issue for humanoid robots arising from inaccurate torque-tracking at the actuator level. Our key idea is to utilize the current feedback from the actuators on the real robot, after training the policy in a simulation environment artificially degraded with poor torque-tracking. Our approach successfully trains a unified, end-to-end policy in simulation that can be deployed on a real HRP-5P humanoid robot to achieve bipedal locomotion. Through ablations, we also show that a feedforward policy architecture combined with targeted dynamics randomization is sufficient for zero-shot *sim2real* success, thus eliminating the need for computationally expensive, memory-based network architectures. Finally, we validate the robustness of the proposed RL policy by comparing its performance against a conventional model-based controller for walking on uneven terrain with the real robot.

INDEX TERMS Bipedal locomotion, humanoid robots, reinforcement learning, *sim2real*.

I. INTRODUCTION

As conventional model-based approaches for humanoid locomotion continue to improve, such as those based on preview control [1] or model predictive control (MPC) [2], their robustness against unexpected disturbances and inaccurate modeling is still an elusive research goal. On the other hand, rapid advancements in RL-based control methods for legged locomotion have shown outstanding performance in unstructured and uncontrolled environments for quadrupedal robots [3], [4], [5], [6] and even bipedal robots [7], [8], [9]. It would be appealing to apply similar methods to develop walking controllers for larger and heavier humanoid robots, too.

Training a capable policy using deep RL is data intensive and can be damaging to the hardware. Physics simulation

The associate editor coordinating the review of this manuscript and approving it for publication was Jiachen Yang^{1b}.

environments offer a safe way to collect a large amount of data, so policies are typically trained in simulation and then transferred to the real system. However, the simulated environment can fail to capture the richness of real-world dynamics. This gives rise to the “reality gap”, more commonly known as the *sim2real* gap. The *sim2real* gap can cause the performance of a policy trained in simulation to drop drastically when deployed on the real hardware. In the case of life-sized humanoid robots such as the HRP-series humanoids, this gap can have a more critical effect on the robot’s stability during walking, compared to the quadrupedal robots or lightweight bipedal robots that are used in most of the recent works. Memory-based policy architectures that can use temporal information to essentially perform online system identification have previously been proposed to tackle this issue [5], [10]. But such networks are generally more computationally expensive than feedforward (FF) networks,



FIGURE 1. HRP-5P humanoid robot trained to perform bipedal locomotion via model-free reinforcement learning in MuJoCo (top); RL policy transferred to the real robot (bottom). We make use of the feedback from measured actuator current to account for the poor torque-tracking on the real system.

which can be prohibitive for prototyping and deployment of RL policies especially for learning tasks with high sample complexity.

Previously, it has been reported that the simulated actuator model has a major impact on the *sim2real* transfer compared to other factors such as link masses and their center-of-mass (CoM) positions [4], [11], [12]. We argue that within the actuator model, the modelling of the torque-tracking behavior forms a key source of discrepancy between the real and simulated robots. In the simulation environment, the joint torque desired by the RL policy is injected exactly into the dynamics computation as the control input. However, in the case of real robots such as the current-controlled HRP-5P, the actual torque exerted by the motor on the link may be significantly different from the desired torque due to imperfect current-control (and a nearly linear relationship between torque and current for brushed DC motors). Since the robot makes environmental interaction only by applying torques on the links, the mismatch between the torque desired by the policy and the actual torque applied on the link may have consequential effects on the control.

In this paper, we develop a system to train bipedal walking policies in simulation and deploy them on the HRP-5P humanoid robot (Figure 1). HRP-5P is a high-power, electrically-actuated, 53 degrees of freedom (DoF) humanoid

robot weighing 105kg, with a height of 182cm [13]. Our key insight is that, on such robots, the *sim2real* gap is significantly a result of imperfect current tracking on the real robot that leads to a mismatch in the desired and applied torque. To this end, we propose to simulate a degraded torque-tracking effect during training and to incorporate current feedback from the motors into the observation space. The resulting policy learns to actively use the current feedback signal and compensate for the inaccurate torque-tracking within the motor drivers. Finally, the policy could be successfully deployed on a real HRP-5P robot.

Specifically, our contributions in this work are as follows:

- By using the proposed *sim2real* approach, we show one of the first demonstration of an end-to-end policy for a real, current-controlled, humanoid robot to achieve dynamic stability. Our policy can achieve forward walking, stepping and turning in-place, and quite standing, by receiving user commands via a joystick.
- We perform ablation study on feedforward Multi-Layered Perceptrons (MLPs) and Long Short Term Memory (LSTMs) networks to show that it is possible to bridge the “reality gap” without relying on memory-based policy architectures or resorting to unreasonably wide randomization of dynamics parameters. This is necessary for the development of RL policies within reasonable amount of computation resources.
- We validate the robustness of the proposed policy on the real robot and compare the qualitative performance of the RL policy to an open-source model-based walking controller for blind locomotion over small uneven and compliant obstacles.

II. RELATED WORK

A. REINFORCEMENT LEARNING FOR LEGGED ROBOTS

Reinforcement Learning has become a powerful approach for synthesizing controllers for legged robots. Control policies are typically trained in simulation and then transferred to the hardware, i.e., *sim2real*. A large number of such works focus on quadrupedal robots, e.g., ANYmal [4], Laikago [14], A1 [5], Jueying [15] and Mini-Cheetah [3]. There are also successes in applying the same approach for bipedal robots, e.g., on the Cassie [7], [10], [16], Digit [17], [18], [19] and NimbRo-OP2X [20]. DeepWalk [20] demonstrates a single learned policy for a real humanoid robot that can achieve omnidirectional walking on a flat floor through the use of Beta policies [21], albeit noting a need for more sophisticated methods for improved transfer.

For the HRP-5P and JVRC-1 humanoids, end-to-end deep RL policies have been previously been demonstrated for walking on planned footsteps, but only in the simulation environment [22]. The focus of this work, on the other hand, is to solve the *sim2real* issues for achieving real robot demonstrations.

Recently, an impressive demonstration of torque-based deep RL locomotion policy for the real TOCABI humanoid robot has shown the advantage of using torque action-space for improved *sim2real* transfer [23]. TOCABI is a human-sized humanoid [24] that can achieve torque-controlled compliance without the use of joint-level torque sensor. Application of such an approach for the low back-drivable joints of HRP-5P humanoid is a matter of future research.

Again, implementing learned policies on the hardware of a bulky humanoid robot, such as Valkyrie and the HRP-series robots, presents significantly greater challenges than lighter robots, especially due to heightened safety risks.

B. SIM2REAL APPROACHES

Important ingredients for successful *sim2real* include (a) careful system identification, e.g., a learned actuator model is incorporated in the simulator to account for hard-to-model actuator dynamics [4], [11], (b) domain adaptation, where the policy learns to adapt based on the history of observations and actions [5], [6], [10], (c) dynamics randomization, where simulation parameters such as mass, friction, inertia, CoM positions are randomized to improve robustness of the policy [25]. While successful on several legged robot platforms, we are yet to see such approaches being successfully applied to life-size (e.g., more than 170cm) humanoid robots with heavy limbs. This is not due to lack of attempt, prior work has explored how to apply reinforcement learning to the NASA Valkyrie robot, e.g., [26], but so far no *sim2real* is demonstrated.

For the case of HRP-5P, end-to-end supervised learning of actuator dynamics [4] is not readily applicable due to absence of joint torque sensors which could provide the ground-truth signal for such an actuator model. Further, search-based system identification methods using real data [11], [12] may be difficult for a large robot as they require multiple trial runs (5)-(10) on the real robot. Domain adaptation and online system identification methods through the use of memory-based agents have been demonstrated successfully for biped robots [5], [10]. However, training memory-based networks (such as LSTMs) can be significantly more computationally expensive and may prohibit real-time inference on the real robot. Our experiments show that with the combination of targeted dynamics randomization, we can bridge the “reality gap” using light-weight FF networks (experimental results in subsection V-F).

C. CONVENTIONAL CONTROL FOR HUMANOID LOCOMOTION

Conventional model-based approaches for humanoid bipedal locomotion consist of local feedback controllers to track Zero-Moment-Point (ZMP) or Center-of-Mass (CoM) trajectories precomputed in an offline process. Stabilization control through the use of divergent component of motion (DCM) has been extensively studied in prior works [27], [28] and

applied to real robots. It relies on the linear inverted pendulum model [29], [30] of bipedal walking. Other recent works such as [1], on the other hand, do not rely on biped-specific dynamics and instead perform online generation of centroidal trajectory based on preview control for impressive multi-contact motion. The method uses a preview control scheme to generate a centroidal trajectory and a stabilization scheme to correct errors in tracking the trajectory. Since the trajectory is generated online the robot can react robustly to environmental disturbances.

Generally, the output of such controllers is in the form of desired ZMP, CoM, and/or contact wrenches, which are then fed to Quadratic programming (QP) solver to compute desired joint positions. The desired joint positions are then tracked by a local proportional-derivative (PD) controller for stiff position control. This is in sharp contrast to RL approaches mentioned above that use low PD gains to achieve compliant tracking, and consequently, offer greater robustness to uneven terrain.

To this end, we perform real robot experiments to compare the robustness of our proposed RL policy to a commonly used, open-source model-based controller for locomotion over uneven terrain (details and results provided in subsection V-E).

III. BACKGROUND

The motor control system on the HRP-5P robot (and generally, on most robot transmission systems) is shown in Figure 2. The block structure consists of a PD controller which computes the desired torque command given the reference position from a higher-level controller and measured position from the joint encoder. The desired torque command — or equivalently, the current command (assuming a proportional relationship between torque and current for a brushed-DC motor) is then sent to a proportional-integral (PI) controller. The PI controller tries to track the current command given the measured current from the current sensor in the motor. The output of the PI controller is fed to the motor power amplifier, which in turn drives the motor.

The key observation here is that the PI controller is unable to precisely track the torque commands, as desired by the higher-level controller or policy, often leading to significant torque-tracking errors. We suspect that such tracking errors could be caused largely due to the effect of the back electro-motive force (EMF). When the motor rotates, the back-EMF creates a counter-voltage that opposes the applied voltage, reducing the current flowing through the armature, leading to tracking errors. This makes the real system vastly different from simulation environments where the desired torque command is applied *exactly* to the actuator without errors. Other factors such as the battery voltage, resistance of the transmission cables, changes in load, or poorly tuned gains of the PI controller may also contribute to poor current tracking.

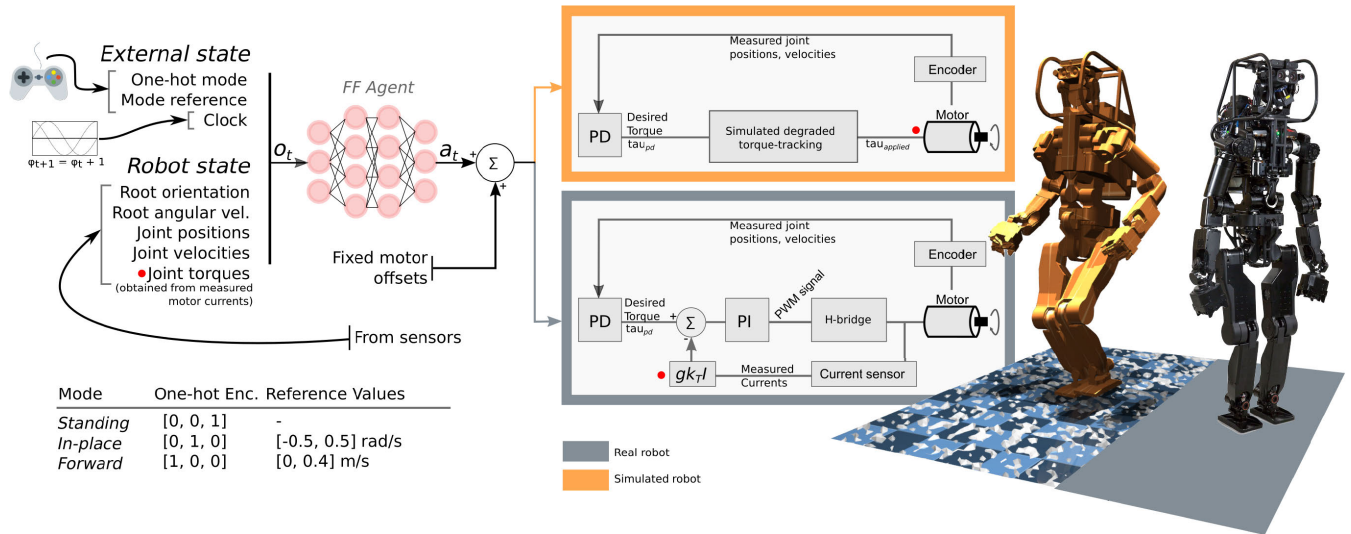


FIGURE 2. Overview of the proposed RL policy with block-diagram of motor control system in HRP-5P. Output of the RL policy in terms of “desired joint positions” is summed with fixed motor-offsets (corresponding to the robot’s nominal pose) and is fed to the joint PD controller. On the real robot, the torque computed by the PD controller is tracked by a PI current controller, albeit, with significant tracking errors. These tracking errors form a crucial component of the *sim2real* gap. To overcome this issue, we propose to (a) simulate a degraded torque-tracking environment during training and (b) observe the applied motor torque, or equivalently, measured motor currents on the real robot (denoted by ●).

IV. APPROACH

In this section, we detail each component involved in training the RL policy (see Figure 2). The training is performed in the MuJoCo simulator [31]. In particular, we describe how we overcome the poor torque-tracking on the real robot by simulating Back-EMF effect and using current feedback for the policy. Since the current and torque on the robot’s actuators are assumed to be proportional, we use the terms interchangeably through the paper.

A. OBSERVATIONS AND ACTIONS

The input to our policy consists of the robot state, the external state, and a clock signal. The robot state includes the joint positions and joint velocities of each actuated joint (6 in each leg), roll and pitch orientation and angular velocity of the root (pelvis), similar to several related works [10], [22]. In this work, we propose to also include the motor-level torque signal for each actuator in the robot state τ_{obs} . In simulation, this signal is equal to the actual torque $\tau_{applied}$ applied to the actuators in the previous timestep. On the real robot, τ_{obs} needs to be derived from the raw current measurements, as explained in V-B. The external state vector comprises of a 3D one-hot encoding to denote the walking mode — [0, 0, 1] for standing and [0, 1, 0] for stepping in-place and [1, 0, 0] for walking forward. It also includes a 1D scalar which acts as a reference value depending on the mode: If the active mode is *Stepping*, the reference value denotes the turning speed; for *Walking* it denotes forward walking speed; and is ignored for the *Standing* mode.

The policy also observes a clock signal that depends on a cyclic phase variable ϕ . This variable is also used to define a periodic reward term in our reward function to generate

walking behaviors. We do a bijective projection of ϕ to a 2D unit cycle:

$$\text{Clock} = \left\{ \sin\left(\frac{2\pi\phi}{L}\right), \cos\left(\frac{2\pi\phi}{L}\right) \right\}, \quad (1)$$

where L is the cycle period. ϕ increments from 0 to 1 at each control timestep and reset to 0 after every L timesteps. Clock is then used as input to the policy.

The policy outputs desired positions of the actuated joints in the robot’s legs. These predictions from the network are added to fixed motor offsets corresponding to the robot’s half-sitting posture. The desired positions are tracked using a low-gain PD controller, which computes the desired joint torque as follows:

$$\tau_{pd} = K_p(q_{des} - q) + K_d(0 - \dot{q}), \quad (2)$$

where K_p and K_d denote the proportional and derivate gain factors respectively. q_{des} is the policy prediction summed with the fixed motor offsets. q and \dot{q} denote the current joint position and velocity.

B. REWARD FUNCTION

Our reward design ensures that a reference motion is not needed. Instead, we rely on several hand-crafted reward terms that promote the desired robot behavior in 3 modes: stand in place, step in-place (including turning) and walk forward given a reference speed. This requires the robot to develop a periodic bipedal gait, follow the mode and reference velocity command and maintain a fixed height. Further, we introduce terms to develop a more realistic motion that will allow *sim2real* transfer in a realistic and safe manner.

1) BIPEDAL WALKING

We introduce reward terms for promoting a symmetrical bipedal walking gait characterized by a periodic motion of the legs, alternating between double-support (DS) phases, and the single-support (SS) phases. Depending on the phase variable ϕ and the desired mode (*standing* or *walking*), the rewards for feet ground reaction forces (GRF) and body speeds are computed.

For example, when ϕ lies in the first single-support region of the gait cycle (the left foot is in the swing phase and the right foot is in the support phase), larger values of GRF on the left foot are rewarded negatively while larger values of GRF on right foot lead to positive reward. Simultaneously, higher speeds for the left foot are incentivized but penalized for the right foot.

The definition of the bipedal walking terms — ground reaction forces at the feet r_{grf} and the feet body speeds r_{spd} — are adopted exactly from [22]:

$$r_{\text{grf}} = I_{\text{left}}^{\text{grf}}(\phi) \cdot F_{\text{left}} + I_{\text{right}}^{\text{grf}}(\phi) \cdot F_{\text{right}} \quad (3)$$

$$r_{\text{spd}} = I_{\text{left}}^{\text{spd}}(\phi) \cdot S_{\text{left}} + I_{\text{right}}^{\text{spd}}(\phi) \cdot S_{\text{right}} \quad (4)$$

where F_{left} and F_{right} denote the GRF and S_{left} and S_{right} denote the body speeds on the left and right foot respectively. We refer the reader to [22] for a detailed explanation of the “phase indicator” functions I_*^{grf} and I_*^{spd} for modulating the reward coefficients for ground reaction forces and feet speeds.

For *standing*, the DS phase is expanded to span the entire gait cycle, and the policy is rewarded to maximize ground reaction forces on both feet while minimizing the feet speeds.

2) ROOT VELOCITY, ORIENTATION AND HEIGHT

The root linear velocity reward term is a simple cost on global speed ${}^x v_{\text{root}}$ of the root link of the robot in the x -direction.

$$r_{\text{rv}} = \exp(-10 \cdot \|{}^x v_{\text{root}} - {}^x \hat{v}_{\text{root}}\|^2) \quad (5)$$

The root yaw velocity term encourages the angular velocity of the root ω_{root} to be close to the desired velocity $\hat{\omega}_{\text{root}}$.

$$r_{\text{av}} = \exp(-10 \cdot \|\omega_{\text{root}} - \hat{\omega}_{\text{root}}\|^2) \quad (6)$$

During training, the active mode is randomly selected between *Standing*, *Stepping*, and *Walking* at the start of an episode. Depending on the active mode, the scalar input for the reference value is sampled from a uniform distribution, i.e., ${}^x \hat{v}_{\text{root}}$ from a range of $[0, 0.4]\text{m s}^{-1}$ if mode is *Walking* and $\hat{\omega}_{\text{root}}$ from a range of $[-0.5, 0.5]\text{rad s}^{-1}$ if mode is *Standing*.

We also reward the policy to maintain the root height h_{root} at a desired value $\hat{h}_{\text{root}} = 0.79$ m:

$$r_{\text{height}} = \exp(-40 \cdot (h_{\text{root}} - \hat{h}_{\text{root}})^2). \quad (7)$$

3) SAFE AND REALISTIC MOTION

In addition to the above terms, we also try to create a motion that remains close to the nominal posture of the robot to avoid

unnecessary sways. This is critical for safe deployment on the real robot, which has a wide range of motion and significantly strong actuators.

To encourage the robot to maintain an upright posture, we use a reward term to minimize the distance between the floor projection of the head position ${}^{x,y} \mathbf{p}_{\text{head}}$ and the root position ${}^{x,y} \mathbf{p}_{\text{root}}$. This prevents the robot from developing a leaned-back behavior and excessively swaying the upper body:

$$r_{\text{upper}} = \exp(-10 \cdot \|{}^{x,y} \mathbf{p}_{\text{head}} - {}^{x,y} \mathbf{p}_{\text{root}}\|^2). \quad (8)$$

We use a term to penalize the distance of the current joint positions \mathbf{q} from the nominal “half-sitting posture”, $\mathbf{q}_{\text{nominal}}$:

$$r_{\text{posture}} = \exp(-\|\mathbf{q} - \mathbf{q}_{\text{nominal}}\|^2). \quad (9)$$

We also place a penalty on joint velocities $\dot{\mathbf{q}}$ above 50% of the maximum joint velocity \mathbf{q}_{lim} .

$$r_{\text{jv}} = \exp\left(-5 \times 10^{-6} \sum_{\dot{\mathbf{q}} > k \cdot \mathbf{q}_{\text{lim}}} \|\dot{\mathbf{q}}\|^2\right). \quad (10)$$

The full reward function is given by:

$$\begin{aligned} r = & w_1 r_{\text{grf}} + w_2 r_{\text{spd}} + w_3 r_{\text{rv}} \\ & + \dots w_4 r_{\text{av}} + w_5 r_{\text{height}} + w_6 r_{\text{upper}} \\ & + \dots w_7 r_{\text{posture}} + w_8 r_{\text{jv}}, \end{aligned} \quad (11)$$

where, the weights $w_1, w_2, w_3, w_4, w_5, w_6, w_7, w_8$ are set to 0.225, 0.225, 0.100, 0.100, 0.50, 0.100, 0.100, 0.100, respectively.

C. DYNAMICS RANDOMIZATION

Since policies trained in simulation interact with an imperfect representation of the world, they are prone to overfitting and show brittleness on the hardware. A common approach to overcome this is to randomize various robot model and environment parameters, such as mass, inertia, motor strength, latency, ground friction, etc [5], [10].

In our work, we carefully select the variables that are needed to be randomized for a better transfer. Firstly, we can expect the mass and position of the center of mass (CoM) of each link to be different on the real system than in the simulation, due to the distribution of electronics and mechanical parts within a link. We randomize the mass of each link by 5% and randomize the CoM positions by 5cm at the start of each episode during training.

Secondly, prior work shows there exists a significant amount of friction between the motor and the load [32] in geared transmission systems. This frictional torque is difficult to identify or even model in simulation. MuJoCo allows the simulation of static friction and viscous friction. Hence, we randomize the static friction magnitude in the uniform range of (2, 8)Nm and coefficient for viscous friction in the uniform range of (0.5, 5)Nm/rad s⁻¹, based on coarse identification performed in [32].

Besides the mass, CoM positions, joint friction, we do not randomize any other robot dynamics parameters during training.

D. TERRAIN RANDOMIZATION

In order to enable the real robot to walk robustly over uneven terrain, we expose the policy to randomized terrains during training. In MuJoCo, the terrains are represented using a height fields - a 2D matrix comprising of elevation data. As generating new height fields online during the training may cause slowness during training, instead, we generate one random height field at the start of the training and then randomize its relative position to the flat floor in each episode. The flat ground plane and the height map are simulated simultaneously, so that the floor resembles a terrain with obstacles of varying heights scattered randomly. In this way the robot is exposed to unevenness of maximum height of 3.5 cm. The introduction of terrain randomization in this work is done uniformly through the training; but it could also be done gradually according to a curriculum to promote smoother learning.

E. SIMULATING BACK-EMF

In order to simulate the phenomenon of poor torque-tracking (observed on the real system), during the training phase, we introduce a modification to the applied torque for each joint at each simulation timestep. The modification is implemented by injecting a counter-torque that scales with the joint velocity, specifically, by using the following equation:

$$\tau_{applied} = \tau_{pd} - k_{bemf} \cdot \dot{q}, \quad (12)$$

where, τ_{pd} denotes the torque at the output of the PD controller, and \dot{q} represents the joint velocity. The damping coefficient, k_{bemf} , is unknown for the real system. During training, we randomize k_{bemf} to simulate different tracking behavior. The coefficient for each joint is sampled uniformly within [5, 40] at every 100 ms. This range was determined empirically by comparing the simulation environment and logs from a real robot experiment.

During training, $\tau_{applied}$ is observed directly by the policy, i.e., $\tau_{obs} = \tau_{applied}$.

V. EXPERIMENTS

A. RL POLICY

1) TRAINING DETAILS

As in [22], both the actor and critic policies are represented by MLP architectures to parameterize the policy and the value functions and use Proximal Policy Optimization (PPO) [33] for training. Both MLP networks have 2 hidden layers of size 256 each and use *ReLU* activations. Each episode rollout spans a maximum of 400 control timesteps (equivalent to 10 s of simulated time), and may reset if a terminal condition is met. Each training batch holds 64 of such rollouts. The learning rate was set to 0.0001. We use the *LOSS* method [8], which adds an auxiliary loss term (in addition to the original

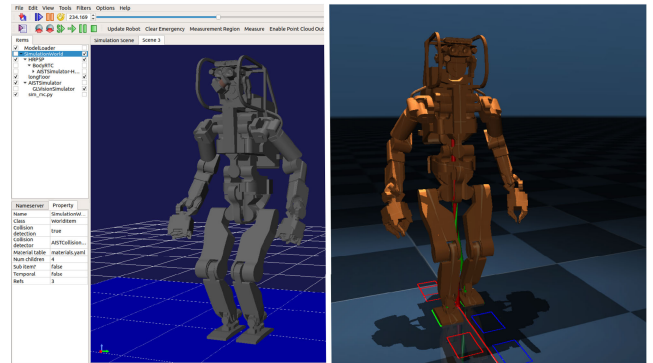


FIGURE 3. Sim-to-sim validation. Simulating HRP-5P in Choreonoid (left) and MuJoCo (right) using the mc-openrtm and mc-mujoco interfaces respectively.

PPO loss term) to enforce symmetry. Training the FF policy takes around 12 hours to collect a total of 120 million samples for learning all modes, on a AMD Ryzen Threadripper PRO 5975WX CPU with 32 cores.

B. IMPLEMENTATION ON REAL ROBOT

We propose to include the actual applied torque in the observation vector to our RL policy. As mentioned previously, the applied torque on the real robot is extracted from the measurements of the current sensors in the motor drivers. The measured current I_{meas} is multiplied by the torque constant k_T and the gear ratio g corresponding to the joint and fed to the policy, as follows:

$$\tau_{mot} = I_{meas} \cdot k_T \quad (13)$$

$$\tau_{obs} = g \cdot \tau_{mot} \quad (14)$$

It is important to note that the applied torque here refers to the torque applied at the level of the actuator scaled to the joint space. The torque applied to the load (i.e. the robot link) cannot be measured on the HRP-5P robot due to absence of joint torque sensor. The difference between the *scaled* actuator torque and the joint torque can in fact be quite significant due to the presence of static friction and viscous friction.

The policy is executed on the control PC of the robot (specifications: Intel NUC5i7RYH i7-5557U CPU with 2 cores, Ubuntu 18.04 LTS PREEMPT-RT kernel), and is implemented as an *mc-rtc*,¹ controller in C++. The inference is done at 40 Hz with the PD controller running at 1000 Hz. Policy inference is quite fast, taking only around 0.2 ms. The global run of the controller is around 1 ms.

C. SIM-TO-SIM VALIDATION

While we use MuJoCo as the training environment, we perform thorough evaluations also in the Choreonoid simulator before real robot deployment. Choreonoid is traditionally a more popular choice for simulating humanoid robot controllers [34]. We use the *mc-rtc* control framework for executing the policy onboard the control PC of the robot.

¹https://jrl-umi3218.github.io/mc_rtc/index.html




FIGURE 4. Baseline (top) v. Proposed policy (bottom): Real experiment logs for torque-tracking on the right hip roll joint. We observe that the torque-tracking becomes unstable for the baseline policy (circled) when the robot is turning and an audible noise can be heard from the joint. Such effects are not observed with the proposed approach, and the robot’s swing leg behaviour visually appears to be more stable.

This allows us to evaluate the same controller code transparently in 3 different environment: (1) in MuJoCo using the *mc-mujoco* interface [35], (2) in Choreonoid and (3) on the real robot, both using the *mc-openrtm* interface for communicating with OpenRTM middleware [36] used in the HRP robots and Choreonoid simulator.

We found the contact modelling in Choreonoid to be more stable than the contact modelling for height fields in MuJoCo. Hence, Choreonoid forms an important part of our pipeline for evaluating policies on uneven terrain before real robot deployment. Nevertheless, besides some small differences, we did not observe any major discrepancies in the policies behavior between MuJoCo and Choreonoid during evaluation.

D. ABLATION STUDY

We perform ablation on the two main ingredients proposed in this work for *sim2real* success - (1) training with simulated poor torque-tracking and (2) providing torque feedback to the policy.

Although real robot experiments are expensive and testing of policies that are prone to failure can be dangerous, developing a test environment in simulation is not a credible alternative. We found that policies that succeed in simulation even in very challenging circumstances (like degraded torque-tracking, uneven terrain, external perturbations), can

still behave undesirably when deployed on the real robot; indicative of a large and critical reality gap. Hence, it is important to study the behaviour of the policies on the real system. We train 3 different policies to analyze the impact of the proposed approach:

- 1) **Policy A** forms the baseline policy. It is trained without simulated poor torque-tracking and without observing the current feedback from the actuators. When deployed on the real robot, this policy gives the worst performance. The robot is prone to self-collision between the feet when the swing leg lands on the ground. This points to the difficulty faced by the policy in controlling the real robot’s leg swing motion. This is because the policy is trained with perfect tracking in the simulation environment but is exposed to degraded tracking on the real robot. Further, we attempt to train another policy with an additional termination condition on the feet distance ($d < 0.2$ m) to promote a wider stance. In this case, self-collision is prevented on the real robot but we can observe that the torque-tracking becomes unstable in some regions of the motion with an audible noise heard from the joints (see Figure 4).
- 2) **Policy B** is trained with poor torque-tracking but without torque feedback. We replace the inputs corresponding to the torque-feedback with the $\mathbf{0}$ vector during training and evaluation, while keeping all other parameters the same. This policy appears robust in the simulation environment. However, when deployed on the real robot, we again observe the self-collision between the feet. The speed of the swing leg is also considerably higher, meaning that the policy finds it harder to compensate for the changing discrepancy between command and applied torque. From this observation, we conclude that providing the torque feedback (from the measured current) is vital for the policy to adapt to the degraded torque-tracking environment.
- 3) **Policy C** is trained using the proposed approach of simulating poor torque-tracking plus providing feedback from current measurements to the policy. This policy appears significantly more stable on the real robot. Self-collision is not observed and there is no audible noise during any part of the motion. The robot could successfully walk upto several meters, including turning, stepping in-place and standing. The robot could also walk over uneven terrain consisting of rigid and soft obstacles upto 2 cm high.

We further analyze the real robot experiment logs corresponding to the 3 policies in Figure 5 for torque-tracking on the “RKP” (right knee pitch) joint. The “RKP” joint is chosen because the tracking errors are more noticeable for this joint. And also because the knee joint is expected to have a more consequential impact on the walking behavior (than compared to the hip yaw joint, say). We observed self-collision between the feet in case of Policy A and Policy B, while Policy C can perform stable walking and

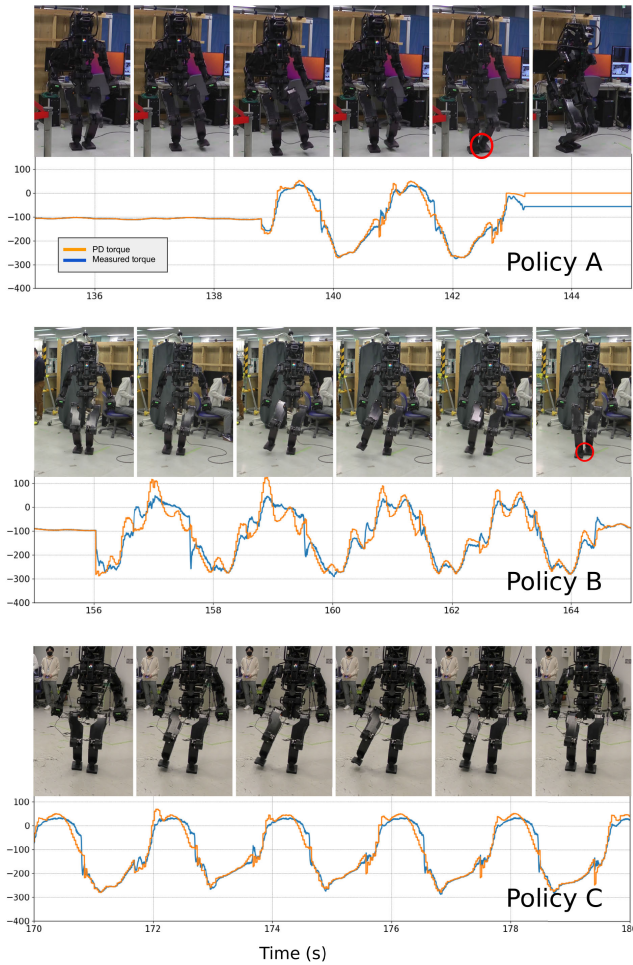


FIGURE 5. Torque-tracking performance on the RKP joint for 3 policies. Policy A is trained without back-emf effect in sim and without feedback. Policy B is trained with back-emf but without feedback. Policy C is trained with back-emf and with the torque (current) feedback. Policy A and B both lead to self-collision between the feet (circled in red) while C is able to perform stable walking.

handle uneven terrain well. The tracking for the “RKP” joint for A and C looks somewhat similar, however, in case of Policy C, the feedback allows the policy to react to the tracking error in the previous timestep and achieve better control on the swing motion. For B, the tracking is observed to be much worse. Tracking for “RCR” (right hip roll) joint for Policy A (retrained for wider feet distance) and C are also shown in Figure 4.

Notably, providing the torque feedback will not eliminate tracking errors because it is difficult for the policy to anticipate the errors in the future. We believe that incorporating history data in the observation space will also be beneficial.

E. COMPARISON TO MODEL-BASED CONTROLLER

We compare the robustness of our policy for locomotion on uneven terrain to an existing model-based controller for humanoid locomotion. For the test, we use the open-source *BaselineWalkingController*,² which provides an

²<https://github.com/isri-aist/BaselineWalkingController>

implementation of walking control based on linear inverted pendulum mode (LIPM). This method combines online CoM trajectory generation based on preview control, ZMP feedback based on the divergent component of motion (DCM) of LIPM, and foot contact force control based on the damping control law [28], [29], [30].

Our test environment consists of a stack of padded carpet tiles, each of thickness 0.6 cm, placed on a flat floor. The robot starts from some distance ahead of the stack and needs to go across while stepping on top of the stack. Since the tiles are made from a soft material, the obstacle forms a compliant support surface - which is more challenging from a balance perspective. The *BaselineWalkingController* could succeed on a stack of 3 tiles but failed on a stack of 4 tiles (nearly 2.4 cm in height) — the robot loses balance and falls when the support leg is on the stacked carpets. On the other hand, the RL policy trained with our approach could succeed on a stack of 5 carpets (= 3 cm high) on 2 of 2 trials. The policy also succeeds in making several partial contacts on the obstacle, where the foot is placed partially on the stack. (The tests are shown in the supplementary video.)

While there exist newer model-based approaches for humanoid locomotion that may provide greater robustness [1], [2], our test still provides valuable insights into the robustness of model-free RL policies against LIPM-based bipedal locomotion controllers. The critical factors responsible for the higher robustness in the case of controllers based on deep RL is subject to debate, but we believe the low PD gains, feedback nature of the policy, and the absence of strict constraints on feet trajectories, may play an important role.

F. FEEDFORWARD vs. MEMORY-BASED POLICIES

In this subsection, we study the behavior of feedforward (FF) MLP policies, without any information about historical states, to that of memory-based policy architectures like Long Short-Term Memory (LSTM) and FF policies with observation history using simulation and real robot experiments. Specifically, we analyzed the following 3 types of network architectures:

- FF policy (vanilla), consisting of a MLP observing only the current observed state o_t
- FF policy with history, which observes the current state along the previous 3 states $\{o_t, o_{t-1}, o_{t-2}, o_{t-3}\}$
- LSTM policy with two hidden layers of 128 units each.

Both FF policies have 2 hidden layers of size 256 nodes. The critic network for each policy had the same architecture as the agent, except for the output layer. We trained each architecture, with and without dynamics randomization (DR). None of the policies were provided any information about the dynamics disturbances. All experiments were performed on a flat and rigid floor. The reward curves are plotted in Figure 6.

During training, we found the FF policies to slightly outperform the LSTM policies in terms of reward collection. This is true even when policies were trained without dynamics randomization, but unsurprisingly, all policies appeared

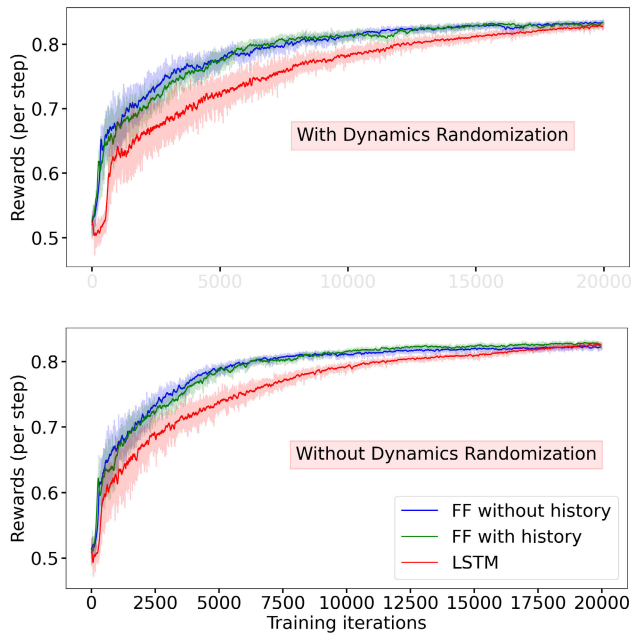


FIGURE 6. Comparison of Training Rewards for different policy architectures. Top: All policies trained with dynamics randomization. FF policy without history (labeled FF (vanilla)) and FF policy with observation history were found to have similar reward curves. LSTM policy achieved the lowest rewards in the beginning, but appeared to converge roughly to the same level. Both FF policies were tested on the real robot, and FF (vanilla) was found to behave most favourably. Bottom: Training without dynamics randomization leads to faster convergence and, again, LSTM achieves the lowest reward. Policies without dynamics randomization were not deployed on the real robot.

to converge faster with no DR. Vanilla FF policy and FF policy with history had nearly overlapping rewards curves. We stopped the training after 20000 iterations, but it is possible that the LSTM reward may eventually surpass FF policies due to overfitting to the simulation dynamics.

FF policies trained without DR were not deployed on the real robot after initial unfavorable results with prototype policies. The swing leg motion visually appeared to be different from the simulated robot and the robot was prone to slight instability. However, we did not observe shakiness or extreme instability even with no DR. With DR, the FF policy with observation history seemed to perform similar to the vanilla FF policy for stepping in-place on the real robot. However, while walking forward and turning we noticed the former to suffer from gradual degradation — gait symmetry appears to break and body oscillations appear to grow. We hypothesize that this is due to the accumulation of *sim2real* errors with time, that would otherwise not occur if the network only observed the current state o_t . We did not deploy the LSTM policy on the hardware as its slower inference speeds may prohibit real time execution. We expect it behave similar to FF policy with history. The vanilla FF policy trained with DR performed the best in all real robot experiments, and hence, emerged as our preferred choice for policy architecture.

Our result appear to be somewhat in contrast to the prior findings for Cassie [10] that report effectiveness of

memory-based networks for achieving *sim2real* success. We believe this is due to the narrow range of randomization performed during training for our robot, meaning that the state-history compression ability of memory-based agents provided no added benefit. Since training the LSTMs takes nearly 3 times as long as the FF policies (36 hours compared to 12 hours for FF), our results show the potential advantage of using FF policies with targeted dynamics randomization for *sim2real* methods. We also note the importance of an accurate initial model of our robot for the same.

VI. DISCUSSION AND CONCLUSION

In this work, we developed a system to train control policies for a life-sized humanoid robot HRP-5P. We identified that the main *sim2real* gap for these types of large robots arises from poor torque-tracking of the motor control systems due to high gear-ratio. We simulated back-EMF and applied torque feedback to the policy to combat the *sim2real* gap. Policies were trained in simulation and directly transferred to the hardware.

Our experiments show that providing the current feedback is a key ingredient for reliable *sim2real* transfer. Without the proposed feedback signal, the policy is prone to failure in controlling the leg swing motion, often causing self-collision between the legs. We could not achieve *sim2real* success without simulating poor torque-tracking during training. For robots with joint-level torque sensors, we believe our proposed approach can yield better performance by accounting for the frictional torque in the joints.

We also provide ablation analysis on the need for memory-based policy architectures. Our results show that a feedforward MLP could be sufficient for successful transfer of policies learned in simulation, in cases where dynamics randomization is performed in a narrow range. Since memory-based networks like LSTMs can be more computationally expensive to train and are prone to overfitting, for this work, we chose MLP policies for real robot deployment.

We could achieve zero-shot transfer without performing aggressive manual tuning of the reward function or randomizing dynamics variables to wide, unreasonable ranges (an often omitted part from the literature). It points to the potential effectiveness of an accurate robot model for training as well as careful identification of key factors responsible in overcoming the *reality gap*. However, the policy appears to make large swaying motion of the swing leg, which could be eliminated by reducing the PD gains of the hip roll joints and relaxing the coefficient of the penalty on joint velocity reward term.

We compared the RL policy to a conventional model-based approach for bipedal locomotion on the real humanoid platform and obtained encouraging results. The RL policy could handle obstacles over 3 cm while the robot lost balance and falls with the model-based controller for obstacles over 2 cm. Although there are newer model-based methods that can tackle larger obstacles, our tests provide promising evidence in the favor of compliant joint tracking, closed-loop

control and relaxed trajectory constraints — offered by the RL approach. Visually, the robot exhibits a cleaner and efficient gait under the model-based controller as opposed to an excessive sway under the RL policy. In the future, this could be tackled through the use of reference motions and fine-tuning of learning parameters.

We release the source code for RL training and evaluation in MuJoCo for reproducibility.³

Future Work: In the future, we plan to expand the framework for developing a policy for *backwards* locomotion and tackle even more challenging terrain. We also hope to identify and overcome other factors inhibiting better *sim2real* transfer.


ACKNOWLEDGMENT

The authors would like to thank all the members of JRL for providing their support in conducting robot experiments that were done during the production of this work and also would like to thank Hiroshi Kaminaga, Mitsuharu Morisawa, and Mehdi Benallegue for the many insightful discussions.

REFERENCES

- [1] M. Murooka, M. Morisawa, and F. Kanehiro, "Centroidal trajectory generation and stabilization based on preview control for humanoid multi-contact motion," *IEEE Robot. Autom. Lett.*, vol. 7, no. 3, pp. 8225–8232, Jul. 2022.
- [2] G. Romualdi, S. Dafarra, G. L'Erario, I. Sorrentino, S. Traversaro, and D. Pucci, "Online non-linear centroidal MPC for humanoid robot locomotion with step adjustment," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, May 2022, pp. 10412–10419.
- [3] G. B. Margolis, G. Yang, K. Paigwar, T. Chen, and P. Agrawal, "Rapid locomotion via reinforcement learning," 2022, *arXiv:2205.02824*.
- [4] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter, "Learning agile and dynamic motor skills for legged robots," *Sci. Robot.*, vol. 4, no. 26, Jan. 2019, Art. no. eaau5872.
- [5] A. Kumar, Z. Fu, D. Pathak, and J. Malik, "RMA: Rapid motor adaptation for legged robots," 2021, *arXiv:2107.04034*.
- [6] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning quadrupedal locomotion over challenging terrain," *Sci. Robot.*, vol. 5, no. 47, Oct. 2020, Art. no. eabc5986. [Online]. Available: <https://robotics.sciencemag.org/content/5/47/eabc5986>
- [7] Z. Xie, P. Clary, J. Dao, P. Morais, J. Hurst, and M. van de Panne, "Learning locomotion skills for cassie: Iterative design and sim-to-real," in *Proc. Conf. Robot Learn. (CORL)*, 2019, pp. 317–329.
- [8] J. Siekmann, Y. Godse, A. Fern, and J. Hurst, "Sim-to-real learning of all common bipedal gaits via periodic reward composition," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2021, pp. 7309–7315.
- [9] J. Siekmann, K. Green, J. Warila, A. Fern, and J. Hurst, "Blind bipedal stair traversal via sim-to-real reinforcement learning," 2021, *arXiv:2105.08328*.
- [10] J. Siekmann, S. Valluri, J. Dao, L. Bermillo, H. Duan, A. Fern, and J. Hurst, "Learning memory-based control for human-scale bipedal locomotion," 2020, *arXiv:2006.02402*.
- [11] W. Yu, V. C. V. Kumar, G. Turk, and C. K. Liu, "Sim-to-real transfer for biped locomotion," 2019, *arXiv:1903.01390*.
- [12] S. Masuda and K. Takahashi, "Sim-to-real transfer of compliant bipedal locomotion on torque sensor-less gear-driven humanoid," 2022, *arXiv:2204.03897*.
- [13] K. Kaneko, H. Kaminaga, T. Sakaguchi, S. Kajita, M. Morisawa, I. Kumagai, and F. Kanehiro, "Humanoid robot HRP-5P: An electrically actuated humanoid robot with high-power and wide-range joints," *IEEE Robot. Autom. Lett.*, vol. 4, no. 2, pp. 1431–1438, Apr. 2019.
- [14] X. B. Peng, E. Coumans, T. Zhang, T.-W. Lee, J. Tan, and S. Levine, "Learning agile robotic locomotion skills by imitating animals," 2020, *arXiv:2004.00784*.
- [15] C. Yang, K. Yuan, Q. Zhu, W. Yu, and Z. Li, "Multi-expert learning of adaptive legged locomotion," *Sci. Robot.*, vol. 5, no. 49, Dec. 2020, Art. no. eabb2174.
- [16] A. Kumar, Z. Li, J. Zeng, D. Pathak, K. Sreenath, and J. Malik, "Adapting rapid motor adaptation for bipedal robots," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Oct. 2022, pp. 1161–1168.
- [17] G. A. Castillo, B. Weng, W. Zhang, and A. Hereid, "Reinforcement learning-based cascade motion policy design for robust 3D bipedal locomotion," *IEEE Access*, vol. 10, pp. 20135–20148, 2022.
- [18] L. Krishna, U. A. Mishra, G. A. Castillo, A. Hereid, and S. Kolathaya, "Learning linear policies for robust bipedal locomotion on terrains with varying slopes," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Sep. 2021, pp. 5159–5164.
- [19] I. Radosavovic, T. Xiao, B. Zhang, T. Darrell, J. Malik, and K. Sreenath, "Learning humanoid locomotion with transformers," 2023, *arXiv:2303.03381*.
- [20] D. Rodriguez and S. Behnke, "DeepWalk: Omnidirectional bipedal gait by deep reinforcement learning," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2021, pp. 3033–3039.
- [21] P.-W. Chou, D. Maturana, and S. Scherer, "Improving stochastic policy gradients in continuous control with deep reinforcement learning using the beta distribution," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 834–843.
- [22] R. P. Singh, M. Benallegue, M. Morisawa, R. Cisneros, and F. Kanehiro, "Learning bipedal walking on planned footsteps for humanoid robots," in *Proc. IEEE-RAS 21st Int. Conf. Hum. Robots (Humanoids)*, Nov. 2022, pp. 686–693.
- [23] D. Kim, G. Berseth, M. Schwartz, and J. Park, "Torque-based deep reinforcement learning for task-and-robot agnostic learning on bipedal robots using sim-to-real transfer," 2023, *arXiv:2304.09434*.
- [24] M. Schwartz, J. Sim, J. Ahn, S. Hwang, Y. Lee, and J. Park, "Design of the humanoid robot TOCABI," in *Proc. IEEE-RAS 21st Int. Conf. Hum. Robots (Humanoids)*, Nov. 2022, pp. 322–329.
- [25] 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," 2018, *arXiv:1804.10332*.
- [26] C. Yang, K. Yuan, W. Merkt, T. Komura, S. Vijayakumar, and Z. Li, "Learning whole-body motor skills for humanoids," in *Proc. IEEE-RAS 18th Int. Conf. Hum. Robots (Humanoids)*, Nov. 2018, pp. 270–276.
- [27] J. Engelsberger, C. Ott, and A. Albu-Schäffer, "Three-dimensional bipedal walking control based on divergent component of motion," *IEEE Trans. Robot.*, vol. 31, no. 2, pp. 355–368, Apr. 2015.
- [28] S. Caron, A. Kheddar, and O. Tempier, "Stair climbing stabilization of the HRP-4 humanoid robot using whole-body admittance control," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, May 2019, pp. 277–283.
- [29] S. Kajita, F. Kanehiro, K. Kaneko, K. Fujiwara, K. Harada, K. Yokoi, and H. Hirukawa, "Biped walking pattern generation by using preview control of zero-moment point," in *Proc. IEEE Int. Conf. Robot. Autom.*, vol. 2, Sep. 2003, pp. 1620–1626.
- [30] S. Kajita, M. Morisawa, K. Miura, S. Nakaoka, K. Harada, K. Kaneko, F. Kanehiro, and K. Yokoi, "Biped walking stabilization based on linear inverted pendulum tracking," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Oct. 2010, pp. 4489–4496.
- [31] E. Todorov, T. Erez, and Y. Tassa, "MuJoCo: A physics engine for model-based control," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Oct. 2012, pp. 5026–5033.
- [32] R. Cisneros, M. Benallegue, R. Kikuuwe, M. Morisawa, and F. Kanehiro, "Reliable chattering-free simulation of friction torque in joints presenting high stiction," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Oct. 2020, pp. 6318–6325.
- [33] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," 2017, *arXiv:1707.06347*.
- [34] S. Nakaoka, "Choreonoid: Extensible virtual robot environment built on an integrated GUI framework," in *Proc. IEEE/SICE Int. Symp. Syst. Integr. (SII)*, Dec. 2012, pp. 79–85.
- [35] R. P. Singh, P. Gergondet, and F. Kanehiro, "Mc-mujoco: Simulating articulated robots with FSM controllers in MuJoCo," in *Proc. IEEE/SICE Int. Symp. Syst. Integr. (SII)*, Jan. 2023, pp. 1–5.
- [36] N. Ando, T. Suehiro, K. Kitagaki, T. Kotoku, and W.-K. Yoon, "RT-component object model in RT-middleware—Distributed component middleware for RT (robot technology)," in *Proc. Int. Symp. Comput. Intell. Robot. Autom.*, 2005, pp. 457–462.

³<https://github.com/rohanpsingh/LearningHumanoidWalking>



ROHAN P. SINGH received the master’s degree from the University of Tsukuba, Japan, in April 2021, where he is currently pursuing the Ph.D. degree. He is also with CNRS-AIST JRL (Joint Robotics Laboratory), AIST, as a Research Assistant. He was a Robotics Engineer with JRL (then known as the Humanoid Robotics Group), from 2017 to 2019. His current research interest includes developing reinforcement learning-based locomotion controllers for humanoid robots.



PIERRE GERGONDET received the master’s degree in embedded systems from École Centrale de Paris and the Ph.D. degree from the University of Montpellier. His Ph.D. thesis in CNRS-AIST JRL, UMI3218/CRT, Tsukuba, Japan, where he is currently a research associate working on embedded software for humanoids. His research interest includes humanoid robotics and the usage of brain–computer interfaces to control a humanoid robot.



ZHAOMING XIE received the B.Eng. degree in computer engineering from The Hong Kong University of Science and Technology, in 2014, the M.S. degree in electrical engineering and mathematics from the Georgia Institute of Technology, in 2017, and the Ph.D. degree in computer science from The University of British Columbia, in 2021. He is currently a Postdoctoral Researcher with Stanford University. His research interests include reinforcement learning, legged robot control, and physics-based character animation.



FUMIO KANEHIRO (Member, IEEE) received the B.E., M.E., and Ph.D. degrees in engineering from The University of Tokyo, Japan, in 1994, 1996, and 1999, respectively. He was a Research Fellow with the Japan Society for the Promotion of Science, in 1998 and 1999. In 2000, he joined the Electrotechnical Laboratory, Agency of Industrial Science and Technology, Ministry of Industrial Science and Technology (AIST-MITI), later reorganized as the National Institute of Advanced Industrial Science and Technology (AIST), Tsukuba, Japan. Since April 2007, he has been a Visiting Researcher with LAAS-CNRS for one year and three months. He is currently the Director of CNRS-AIST JRL (Joint Robotics Laboratory), IRL, AIST. His research interests include the software platform development and whole body motion planning of the humanoid robot.

...