

Template Model Inspired Task Space Learning for Robust Bipedal Locomotion

Guillermo A. Castillo¹, Bowen Weng¹, Shumpeng Yang², Wei Zhang², and Ayonga Hereid³

Abstract—This work presents a hierarchical framework for bipedal locomotion that combines a Reinforcement Learning (RL)-based high-level (HL) planner policy for the online generation of task space commands with a model-based low-level (LL) controller to track the desired task space trajectories. Different from traditional end-to-end learning approaches, our HL policy takes insights from the angular momentum-based linear inverted pendulum (ALIP) to carefully design the observation and action spaces of the Markov Decision Process (MDP). This simple yet effective design creates an insightful mapping between a low-dimensional state that effectively captures the complex dynamics of bipedal locomotion and a set of task space outputs that shape the walking gait of the robot. The HL policy is agnostic to the task space LL controller, which increases the flexibility of the design and generalization of the framework to other bipedal robots. This hierarchical design results in a learning-based framework with improved performance, data efficiency, and robustness compared with the ALIP model-based approach and state-of-the-art learning-based frameworks for bipedal locomotion. The proposed hierarchical controller is tested in three different robots, Rabbit, a five-link underactuated planar biped; Walker2D, a seven-link fully-actuated planar biped; and Digit, a 3D humanoid robot with 20 actuated joints. The trained policy naturally learns human-like locomotion behaviors and is able to effectively track a wide range of walking speeds while preserving the robustness and stability of the walking gait even under adversarial conditions.

I. INTRODUCTION

Robust bipedal robot locomotion presents a challenging problem for robotics research due to the complexity of high dimensional models, unilateral ground contacts, and nonlinear and hybrid dynamics. Common methods applied in bipedal locomotion rely on solving optimization problems using the robot's full-order or reduced-order model to find feasible trajectories that realize stable walking gaits. In general, using full-order models results in computationally expensive problems that cannot be solved in real-time [1], [2]. To reduce the computation time, reduced-order models are used to capture the dynamics of the full-order system and plan trajectories for the robot's center of mass (CoM) and end-effectors. However, the assumptions made on reduced-order models such as a constant CoM height limit their performance on dynamic locomotion behaviors and their

This work was supported in part by the National Science Foundation under grant FRR-21441568, and by the National Natural Science Foundation of China under Grant No. 62073159.

¹Electrical and Computer Engineering, Ohio State University, Columbus, OH, USA; {castillomartinez.2,weng.172}@osu.edu.

²SUSTech Institute of Robotics, Southern University of Science and Technology (SUSTech), China; 11930364@mail.sustech.edu.cn, zhangw3@sustech.edu.cn.

³Mechanical and Aerospace Engineering, Ohio State University, Columbus, OH, USA. hereid.1@osu.edu.

accuracy to predict the behavior of the real robot under certain conditions. Recently, the angular momentum-based linear inverted pendulum (ALIP) has been presented as an improved alternative to the Linear Inverted Pendulum (LIP) to predict the evolution of the model's state, demonstrating in simulation and hardware experiments that the angular momentum about the contact point can be more accurately predicted than the CoM velocity [3], [4].

With the recent success of deep learning in tackling challenging control problems, machine learning-based approaches have exploited advances in physics simulators and computing power to learn locomotion policies through more structured learning frameworks. Learning from motion references has become a popular choice to exploit large amounts of data to train walking policies. The data is obtained from motion capture systems, public motion data sets, or even from video clips, and it is used as goal references in the reward design [5]–[7].

Other learning methods rely on using optimization to obtain a single feasible reference trajectory [8], [9], or libraries of reference trajectories [10], [11] to guide the learning. However, these approaches require large amounts of data, and the learned policy often lacks interpretability and control over the parameters of the walking gait. This makes it difficult to adjust the policy during the sim-to-real transition. As an alternative, more complex frameworks have been proposed to combine learning algorithms with model-based controllers. The authors in [12] take insights from the Hybrid Zero Dynamics (HZD) to learn joint trajectories for planar robots. In [13], an HZD-based approach is used to learn a policy that satisfies Control Barrier Functions (CBF) defined on the reduced-order dynamics.

In this work, we propose a hierarchical RL-based approach to address bipedal locomotion in underactuated and fully actuated robots. At the HL stage, RL is used to train policy that learns task space commands for different walking speeds. At the LL, a model-based nonlinear controller is implemented to track the trajectories generated by the HL planning. Several RL-based approaches have been already proposed to exploit hierarchical structures. In [11], [14], a task space policy is trained to walk at different speeds. However, the method relies on solving a series of optimization problems using the Spring Linear Inverted Pendulum to create a gait library that is used as a reference for the reward and the target end-effector positions. The policy learns residual terms that are added to the task space references [14] or joint space references [11]. Different from these approaches, our method directly learns a set of task space actions that completely

characterize the dynamic walking gait without the need for previously computed reference trajectories. Moreover, we use different state and action spaces that significantly simplify the complexity of the learning problem and can be generalized to both unactuated and underactuated robots.

In our previous work [15], a cascade structure is implemented to compensate the learned trajectories with feedback regulators to increase the robustness of the walking gait [16], [17]. Although the method was successfully tested in hardware, the interpretability of the learned policy was limited by the complex structure imposed over the input-output mapping of the RL policy and the addition of compensation terms on top of the learned joint trajectories. On the one hand, the integration of the feedback regulators improves the robustness and sim-to-real transfer of the learned policy. On the other hand, it makes it difficult to identify the actual contribution of the learned policy to the robustness of the walking gait. This results in a policy limited to naturally exploiting the state and action spaces and a restricted walking speed range with the robot Digit, e.g., $v_x \in [-0.5, 0.5]$ m/s.

In this work, we propose a more efficient and clean framework that completely decouples the HL learning policy from the LL controller with better insights into the selection of the state and action spaces that results in improved sample efficiency and interpretability of the policy. We demonstrate the proposed framework is general for 2D and 3D bipedal robots and can be applied even in the case of underactuated robots. Moreover, we show that the learned policy achieves enhanced performance and robustness compared with our previous work [15].

The main contributions of this paper are as follows: 1) **A simple, efficient, and general** hierarchical learning framework that fully decouples the HL planner from the LL feedback controller. Different from other task space learning approaches, our method (i) uses a reduced-order state for the RL, (ii) learns to walk from scratch, and (iii) computes a set of task space actions that fully characterize dynamic walking gaits. The selection of inputs and outputs is general to bipedal robots of different morphology and degrees of freedom. We show results for actuated and underactuated 2D (Rabbit, Walker2D) and 3D robots (Digit).

2) **Insightful design of the RL state space.** We use the ALIP state and speed tracking information to design a reduced-order state space for the RL that captures the complex dynamics of bipedal locomotion while simplifying the learning process.

3) **Enhanced flexibility of the policy** to naturally exploit the nonlinear dynamics of bipedal locomotion. By including the desired step length, torso orientation, and CoM's height in the action space of the RL, the policy is not restricted to particular behaviors. This allows the policy to learn natural behaviors seen in dynamic locomotion without enforcing them during training.

4) **A robust locomotion controller** that accurately tracks a wide range of walking speeds, even under external disturbances and challenging terrains, with inclinations up to 20 degrees for both underactuated and fully-actuated robots.

II. PRELIMINARIES AND PROBLEM FORMULATION

A. Bipedal locomotion as a hierarchical problem

In general, the bipedal locomotion problem can be characterized as a hybrid system determined by a collection of phases of continuous dynamics with discrete events between the transitions of the continuous phases. Formally, the hybrid system model for biped locomotion can be defined as

$$\Sigma : \begin{cases} \dot{x} = f(x) + g(x)u + \omega(x, u) & x \in \mathcal{X} \setminus \mathcal{H} \\ x^+ = \Delta(x^-) & x^- \in \mathcal{H}, \end{cases} \quad (1)$$

where $x \in \mathcal{X} \subseteq \mathbb{R}^n$ denotes the robot states, $u \in \mathcal{U} \subseteq \mathbb{R}^m$ is a vector of actuator inputs, and $\omega \in \Omega \subseteq \mathbb{R}^w$ a vector of disturbances and uncertainties. The switching surface \mathcal{H} is typically the hyper-surface of points corresponding to the height of the swing leg above the ground being zero, and the reset map $\Delta : \mathcal{H} \rightarrow \mathcal{X}$ denotes the post-impact state values x^+ immediately after switching as a function of the pre-impact state values x^- right before switching.

The control of the bipedal locomotion system described by equation (1) can be formulated as a hierarchical control problem composed by a HL planner and a LL tracking controller. This cascade structure is presented in Fig. 1. The high level policy π_y generates trajectories to realize walking gaits according to design parameters and HL commands, e.g., average walking speed, robustness, terrain slope, etc. The LL policy π_m computes the actuator inputs to track the desired trajectories commanded by the HL planner.

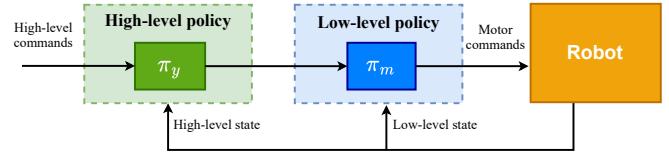


Fig. 1: Hierarchical structure for bipedal locomotion

The general structure presented in Fig. 1 can characterize most controller formulations used in both model-based and model-free state-of-the-art methods for bipedal locomotion. Once the HL trajectories have been generated by the policy, classic model-based control approaches such as feedback linearization, inverse dynamics QP, operational task space controllers, or simple PD controllers can be used to track the desired HL commands. The choice of the LL control policy π_m will mostly depend on the action space of the HL policy, e.g., joint space versus task space.

B. Reduced order models for HL planning

Reduced order models have become a powerful tool for the design of HL planners for bipedal locomotion since they allow using simple dynamical models to characterize biped walking behaviors. Recently, the ALIP model has gained attention because of its advantages over LIP to predict the evolution of the state space. The learning-based approach proposed in this work is heavily inspired by recent results using ALIP as a step planner [3], [4].

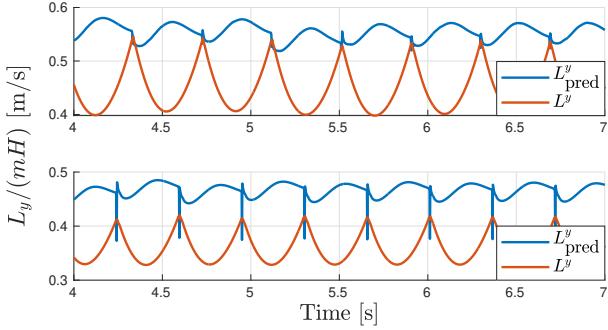


Fig. 2: Prediction of L^y at the end of each step under ideal (top) and non-ideal (bottom) conditions.

Angular Momentum-based Linear Inverted Pendulum (ALIP): Considering the states $\{x, L^y\}$, where x is the CoM position in x direction and L^y is the pitch component of the angular momentum about the contact point, the ALIP dynamics is given by

$$\begin{bmatrix} \dot{x} \\ \dot{L}^y \end{bmatrix} = \begin{bmatrix} 0 & 1/(mH) \\ mg & 0 \end{bmatrix} \begin{bmatrix} x \\ L^y \end{bmatrix}, \quad (2)$$

where m is the total mass and H is the constant CoM height. The main advantage of using the ALIP model over the LIP model is that the evolution of the angular momentum about the contact point is closer to its behavior on the full-order robot's model and the actual hardware [3].

1) *Limitations:* Although ALIP does better work describing the actual behavior of the system than LIP, both are linear models subject to assumptions such as point mass body, constant CoM height, and the angular momentum about the CoM being zero during the walking gait.

In addition, the prediction of the state at the end of the step depends on the step duration T . This implies that an accurate prediction would depend on the perfect timing of the touchdown event, which could only happen in ideal conditions, e.g., perfect tracking of the LL controller, point-contact foot, and non-irregular walking surfaces.

To analyze this effect, we compare the predicted value of L^y scaled by mH at the end of the step with the actual L^y for the five-link bipedal robot Rabbit in Fig. 2. We show the evolution of L^y in simulation using the MuJoCo physics engine [18]. To simulate ideal conditions on the model as closely as possible, we set the geometry of the robot's links to be very thin (to emulate point contact with the ground) and use a LL feedback linearization controller with high gains to encourage better tracking and accurate touchdown timing. For the "non-ideal" conditions, we use the real geometry and dynamic properties of the robot's links (as described in [19]), and we use an inverse dynamics QP controller [20]. The results show that under non-ideal conditions, the prediction of L^y at the end of the step differs significantly from its actual value.

C. Task-space LL controller

Several approaches have been proposed in the literature to design task space controllers that consider the full-order

model of the legged robot. Considering a mechanical system with configuration space \mathcal{Q} and generalized coordinates $q \in \mathcal{Q}$, the equations of motion formulated using the method of Lagrange are given by:

$$M(q)\ddot{q} + H(q, \dot{q}) = Bu + J^T(q)\lambda \quad (3)$$

$$J(q)\ddot{q} + \dot{J}(q, \dot{q})\dot{q} = 0, \quad (4)$$

where $D(q)$ is the inertia matrix, $H(q, \dot{q}) = C(q, \dot{q})\dot{q} + G(q) + F$ is the vector sum of the Coriolis, centripetal, gravitational, and additional non-conservative forces, B is the actuation matrix, and $J(q)$ the Jacobian of the holonomic constraints.

We denote that the system (3) can be expressed in the general form (1). Let $x = (q^T, \dot{q}^T)^T \in T\mathcal{Q} = \mathcal{X}$, then

$$f(x) = \begin{bmatrix} \dot{q} \\ -D^{-1}(q) (J^T(q)\lambda - H(q, \dot{q})) \end{bmatrix} \quad (5)$$

$$g(x) = \begin{bmatrix} 0 \\ D(q)^{-1}B \end{bmatrix}. \quad (6)$$

The task space feedback controller tracks a set of desired trajectories of the form:

$$y(x) = y^a(x) - y^d(\tau(x)), \quad (7)$$

where y^a and y^d are smooth functions, and y^d characterizes the desired behavior of the system. Upon the assumption that $y(x)$ has relative degree 1 or 2, nonlinear control methods can be applied to find a control law that drives $y(x)$ to zero, which implies the outputs converge to their target values.

III. METHOD

This section presents the methodology for the design of the proposed learning-based hierarchical controller for bipedal locomotion. First, we introduce the overall structure of the framework. Then, we describe the learning-based HL and model-based LL components of the framework.

A. Hierarchical structure for bipedal locomotion

The proposed learning-based framework combines the capabilities of model-based and model-free methods into a hierarchical structure to realize robust locomotion controllers for underactuated and fully-actuated bipedal robots. Inspired by the success of reduced-order models for the online generation of HL trajectories, we use reinforcement learning to train an HL policy that maps a reduced state space inspired by the ALIP model to a set of task space commands to generate online task space trajectories for the robot's base and end-effectors. For the LL task space controller, we use well-known model-based inverse dynamics controllers to guarantee the tracking performance of the system's outputs.

The proposed hierarchical structure is presented in Fig. 3. By combining the learning-based HL planner with the model-based LL controller, we obtain a robust controller capable of accurately tracking a wide range of walking speeds while preserving a good tracking performance for the task space trajectories. This significantly increases the flexibility and safety of the policy when compared with pure learning-based controllers.

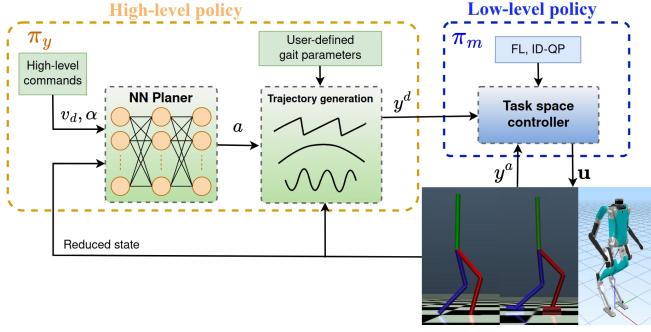


Fig. 3: Overall structure of the proposed learning-based framework. The HL policy maps a reduced-order state to task space trajectories that are tracked by the LL policy.

B. Reinforcement Learning for High-Level Planning

The problem of determining a motion policy for bipedal robots can be modeled as a Markov Decision Process (MDP), which consists of a tuple of components defined as

$$\mathcal{M} := (\mathcal{S}, \mathcal{A}, \mathbb{P}, r, \xi, \gamma). \quad (8)$$

Here \mathcal{S} is the state space, and \mathcal{A} is the set of feasible actions referred to as the action space. Specifically, at time t , an agent (i.e., the motion planner) takes an action $a_t \in \mathcal{A}$ at state $s_t \in \mathcal{S}$, transits into the next state $s_{t+1} \in \mathcal{S}$ according to the transition probability $\mathbb{P}(s_{t+1}|s_t, a_t)$ and receives a reward $r(s_t, a_t, s_{t+1})$. Moreover, ξ denotes the distribution of the initial state $s_0 \in \mathcal{S}$, and $\gamma \in (0, 1)$ denotes the discount factor. The stochastic transition of the MDP process captures the random sampling of initial states in the policy training and dynamics uncertainty due to model mismatch and random interactions with the environment (e.g., early ground impacts).

1) *Reduced-Order State Space*: Several works have already proposed using a reduced state of the robot as the observation space of the learning algorithm. However, the choice of the reduced state is made based on trial and error or empirical observations of the policy performance. In this work, we leverage recent results on the effectiveness of using angular momentum about the contact point to regulate the walking speed of biped robots [3]. Inspired by the ALIP model, we select the state

$$s = (x, y, L^x, L^y, e_{\bar{v}^x}, e_{\bar{v}^y}, v_x^d, v_y^d, \alpha). \quad (9)$$

where (x, y, L^x, L^y) is the ALIP state composed by the robot's base x and y position and the angular momentum about the contact point along the x and y axes, $(e_{\bar{v}^x}, e_{\bar{v}^y})$ is the error between the average velocity the robot's base (\bar{v}_x, \bar{v}_y) and the desired robot's velocity (v_x^d, v_y^d) , and α is the terrain slope measured in radians. We denote that we use the robot's base position instead the CoM because of practical convenience for future hardware experiments. The CoM estimation on complex robots may result in noisy measurements, while the base position with respect to the contact point can be easily computed using forward kinematics.

We assume the slope of the terrain is known by the learning agent. This assumption is reasonable since most of the bipedal robots available for research and commercial applications are equipped with perception systems to map the surrounding environment. Even in the absence of perception systems, proprioceptive approaches could be used to accurately estimate the terrain slope based on the orientation of the robot's base and feet, as we have shown in simulation and hardware in our previous work [15].

2) *Task Action Space*: The action $a \in \mathcal{A}$ is chosen to be

$$a = (p_{sw,T}^x, p_{sw,T}^y, q_\phi, h^d) \quad (10)$$

where $p_{sw,T}^x, p_{sw,T}^y$ correspond to the position of the swing foot w.r.t. the robot's base at the end of the swing phase T , i.e., the landing position of the swing foot, q_ϕ is the absolute torso pitch angle, and h^d is an offset added to the nominal height of the robot's base w.r.t. the stance foot. This selection of the action space encourages the flexibility of the policy to exploit the natural nonlinear dynamics of the biped robot and enhance the robustness of the policy under big disturbances, sudden speed changes, and walking at high speeds, as it will be shown in section V.

The HL actions a are used to generate smooth task space trajectories for the robot's floating base and end-effectors. The trajectory for $p_{sw,T}^{x,y}$ is generated using a minimum jerk trajectory of a straight line segment connecting $p_{sw,0}^{x,y}$ with $p_{sw,T}^{x,y}$. The swing foot position at the beginning of the walking step, $p_{sw,0}^{x,y}$ is computed using Forward Kinematics (FK) and updated at every touch-down event. The desired position for the swing foot at landing, $p_{sw,T}^{x,y}$ is updated by the HL policy at the frequency of 30Hz. The vertical trajectory of the swing foot position w.r.t. the robot's base is generated using a 5th order Bézier Polynomial parameterized by the vertical position of the foot at the beginning ($p_{sw,0}^z$) and the end ($p_{sw,T}^z$) of the step, and the high foot clearance $p_{sw,T/2}^z$. For flat ground terrain, we have

$$p_{sw,T}^z = -h^d \quad (11)$$

We update $p_{sw,T}^z$ by

$$p_{sw,T}^z = -h^d + p_{sw,T}^x * \tan(\alpha) - p_{off}^z, \quad (12)$$

where α is the terrain slope and $p_{off}^z = 0.005\text{m}$ is a small offset added to guarantee the swing foot makes contact with the ground.

The Neural Network chosen to parameterize the HL policy is a Recurrent Neural Network with 2 hidden layers, each layer with 128 units for the case of 2D robots and 256 units for 3D robots. The hidden layers use the ReLU activation function, and the output layer is bounded by the sigmoid activation function and a scaling factor to constrain the maximum value of the HL commands.

C. Low-level task space controller

The LL task space controller is designed using standard techniques of the nonlinear systems control literature. In particular, we implement two types of model-based controllers i) Feedback Linearization (FL), and ii) Inverse Dynamics with

QP formulation (ID-QP). We evaluate the performance of the HL policy with different LL controllers and show that the learned policy is robust to any choice of the LL controller. The purpose of this evaluation is to demonstrate the versatility of the task space-based HL planner to adapt to different LL control structures without affecting the performance of the learned policy. This also provides more flexibility for the designer to use any LL control approach at their convenience. For instance, FL is easy to implement and requires less computation time, but it is known to be hard to implement on real hardware. Therefore, FL could be used during the training process of the HL policy, while any suitable ID-QP formulation could be used for hardware experiments.

For more details, we refer the reader to [20], [21], where several QP formulations for bipedal locomotion are proposed with successful applications to real hardware. In this work, we use the most basic case of the ID-QP formulation in [20] for the 2D robots and the Task Space Inverse Dynamics (TSID) formulation in [20] for the 3D robot Digit.

D. Learning procedure

The reinforcement learning algorithm we use in this work is an implementation of the Proximal Policy Optimization [22] algorithm with parallel experience collection, input normalization, and fixed covariance. The algorithm shares the same code base as the implementations in [11] and [9].

For each episode, the initial state of the robot is set randomly from a normal distribution about an initial pose corresponding to the robot standing in the double support phase. One iteration of the HL policy corresponds to the interaction of the learning agent with the environment. The HL policy takes the reduced-order state $s \in \mathcal{S}$ and computes an action $a \in \mathcal{A}$ that is converted in desired task space trajectories y^d at the time t . The reference trajectories are then sent to the LL task space controller. The LL control loop runs at a frequency of 1 KHz, while the HL planner runs at 33 Hz. The maximum length of each episode is 300 steps, which corresponds to 9 seconds of simulated time.

The episode has an early termination if any of the following conditions are violated:

$$|q_\phi| < 1\text{rad}, \quad h < 0.5\text{m}. \quad (13)$$

The simple reward function (14) adopted in this work is designed to keep track of the target walking speed while realizing a stable walking gait. In particular, the terms r_{v_x}, r_{v_y} encourage the tracking of the longitudinal and lateral target speeds. Since the torso pitch angle is part of the learning action space, the term $r_{L_{\text{CoM}}}$ encourages the policy to avoid excessive changes in the torso orientation without explicitly restricting the torso pitch angle. Finally, the term r_a encourages the policy to avoid excessive variations between the last action and the current action. This avoids unnecessary overshooting in the commanded actions that may produce risky behaviors during the walking gait. The weighted reward function is given as:

$$\mathbf{r} = \mathbf{w}^T [r_{v_x}, r_{v_y}, r_{L_{\text{CoM}}}, r_a]^T, \quad (14)$$

where

$$r_{v_x} = \exp(-\|\bar{v}_x - v_x^d\|^2) \quad (15)$$

$$r_{v_y} = \exp(-\|\bar{v}_y - v_y^d\|^2) \quad (16)$$

$$r_{L_{\text{CoM}}} = \exp(-\|L_{\text{CoM}}\|^2) \quad (17)$$

$$r_a = \exp(-\|a_k - a_{k-1}\|^2). \quad (18)$$

and \mathbf{w}^T is a vector of weights corresponding to each reward term. For 2D robots we use $\mathbf{w}^T = [0.6, 0, 0.2, 0.2]$ while for 3D robots we use $[0.3, 0.3, 0.2, 0.2]$.

IV. ILLUSTRATION EXAMPLE

In this section, we show the proposed method can be generalized to both underactuated and fully actuated robots without any changes to the structure of the HL planner policy. Moreover, we demonstrate the framework can be applied in 2D and 3D bipedal robots. We use 3 different robots.

Rabbit is a five-link, planar underactuated bipedal robot with point feet and four actuated joints, two in the hips and two in the knees. Despite its simple mechanical structure, Rabbit still provides a suitable representation of biped locomotion, which is the reason it has been considered as a test bed for advanced control theory in the field of legged robots [19].

Walker2D is a seven-link, planar, fully actuated bipedal robot with 6 actuated joints, two in the hips, two in the knees, and two in the ankles. The additional degrees of freedom at the ankles enable the robot to realize human-like walking gaits and balancing.

Schematics of the Walker2D and Digit are shown in Fig. 4. Rabbit shares the same design and structure as Walker2D without the feet and ankle joints.

Digit is a 3D fully actuated bipedal robot with 30 DoF and 20 actuated joints. Each leg has six actuated joints corresponding to the motors located on the robot's hip, knee, and ankle and two passive joints corresponding to the robot's shin and tarsus joints. In addition, it has four actuated joints per arm corresponding to the shoulder and elbow joints. Fig. 4 shows the kinematic structure of Digit.

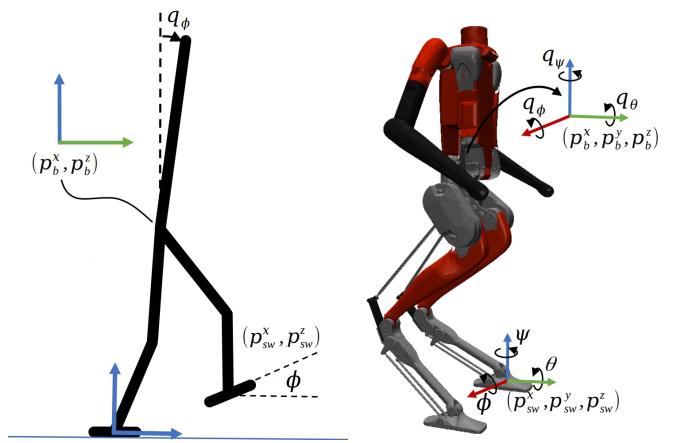


Fig. 4: Schematics of the robots Walker2D and Digit.

A. Task-space outputs for the LL controller.

The set of task space outputs of relative degree 2 described by equation (7) to characterize the walking gait of the biped robot are defined as follows:

$$y_2^a(q) := \begin{bmatrix} q_\phi \\ h \\ p_{sw}^x \\ p_{sw}^y \\ p_{sw}^z \\ \phi_{sw} \end{bmatrix} \rightarrow \begin{pmatrix} \text{torso pitch angle} \\ \text{base height} \\ \text{swing foot } x \\ \text{swing foot } y \\ \text{swing foot } z \\ \text{swing foot pitch} \end{pmatrix} \quad (19)$$

This selection of outputs is common in the field of bipedal locomotion. The first five outputs ($q_\phi, h, p_{sw}^x, p_{sw}^y, p_{sw}^z$) are valid for both underactuated and fully actuated robots. However, for an underactuated robot, it is not possible to control the horizontal position of the robot's base or CoM. Therefore, the evolution of the base velocity is indirectly controlled by the HL learned policy through the planning of touchdown position. In the case of fully actuated robots, we also consider the sixth output to control the swing foot pitch angle to be parallel to the walking surface. This contributes to reducing disturbances at the touchdown event. Although we could add an additional output (ϕ_{sw}) to control the horizontal position and velocity of the robot's base, we prefer to rely on the HL planning to control the robot's speed indirectly. The objectives of this choice of design are twofold: i) devise a general framework for both underactuated and fully-actuated robots that share the same structure for the HL policy independently of the particular design of the bipedal robot. ii) simplify the design of the controller and avoid limitations of the torque ankle to control the robot's base position. Although this effect could not be significant for quasi-static locomotion gaits, it does matter when realizing agile and dynamic locomotion.

V. SIMULATION RESULTS

In this section, we show the performance of the learned HL policy under different testing scenarios with three different robot models, including Rabbit, Walker2D, and Digit. Moreover, we analyze the contribution of the HL policy to the robustness of the walking gait, and we compare our method with similar model-based and model-free approaches.

A. Speed tracking for different velocity profiles.

We test the learned policy for tracking a velocity profile in different directions. Fig. 5 shows the velocity tracking performance of the learned HL policy for the robots Rabbit and Walker2D. To evaluate the robustness of the policy under different LL controllers, we test the same policy with the ID-QP controller and the FL controller with different tracking gains. The results show the policy effectively tracks the walking speeds in the range $[-1, 1]$ m/s, even with aggressive changes in the velocity profile.

To highlight the contribution of the choice of action space, we present in Fig. 6 the actions computed by the HL policy for different commanded velocities. When a steep change in the desired velocity is commanded, the policy

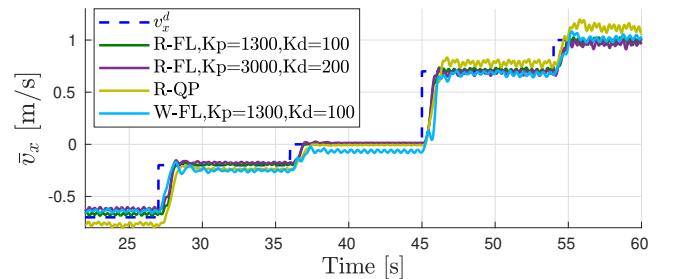


Fig. 5: Velocity tracking performance of the learned policy with different LL controllers. The prefix R/W is used to differentiate a policy for Rabbit or Walker2D.

uses the torso orientation to compensate for variations in the robot's speed and angular momentum. We denote that these behaviors are not enforced during the training process but arise naturally from the insightful design of the proposed framework. Interestingly, some of these strategies are also observed in human locomotion [23].

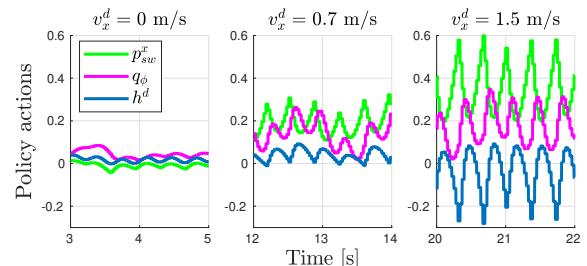


Fig. 6: Contribution of the policy actions for different speeds.

B. Comparison with ALIP model-based approach

To assess the advantages of the proposed HL planner with respect to pure model-based approaches, we compare the performance between our learning-based controller and the ALIP-based controller. Fig. 7 (top) shows the HL-RL policy (ours) outperforms the model-based controller in tracking a velocity profile, especially for high speeds. We also show the variation of L_{CoM} (bottom) to denote the trade-off the policy learns between minimizing L_{CoM} and tracking v_x^d . For small speeds, L_{CoM} looks quite similar for both controllers. For high speeds, the policy learns to prioritize speed tracking over minimizing L_{CoM} . This behavior is encouraged by the selection of weights in the reward function (14).

C. Comparison with other RL-based approaches

To highlight the contribution of the proposed framework in terms of speed tracking for 3D bipedal robots and sample efficiency, we compare our method with our previous work in [15] and the end-to-end learning approach presented in [9]. Although there are no end-to-end learning approaches implemented on the Digit robot in the literature, there are several works that have done so with the robot Cassie, which shares the same leg morphology as Digit. Therefore, we choose to implement the method in [9] with Digit as it focuses on speed tracking, which makes it more comparable to our method. Moreover, the framework in [9] is the base for several SOTA end-to-end learning approaches for bipedal

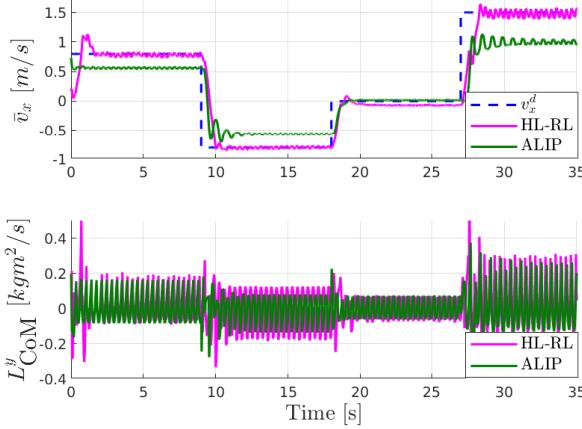


Fig. 7: Comparison of HL-RL with ALIP-based controller.

locomotion [14], [24], [25]. Finally, the code implementation for [9] is publicly available online, which makes the comparison as fair as possible in terms of the reproducibility of their work.

In Fig. 8, we present the comparison results for speed tracking on flat ground with the robot Digit using our learned HL policy, the HZDRL controller in [15], and the end-to-end RL policy in [9]. We observe that the HZDRL controller fails, i.e., the robot falls, for speeds higher than 0.5 m/s, while the end-to-end RL controller fails to track low speeds accurately and realizes a non-smooth walking motion that causes higher variance in the speed profile. This effect may be caused by the reference trajectory used to guide the learning. In our implementation of [9], we use a reference trajectory corresponding to Digit walking forward at a speed of 0.8 m/s. For more details, the reader can refer to [9]. Our learned policy can successfully track the desired walking speed in a significantly wider range compared with the other methods. Additional testing with Digit demonstrates our controller can handle desired walking speeds in the range $v_x \in [-1.0, 1.5]\text{m/s}$ and $v_y \in [-0.5, 0.5]\text{m/s}$, including combinations of both, i.e., diagonal walking, as can be seen in the accompanying video submission.

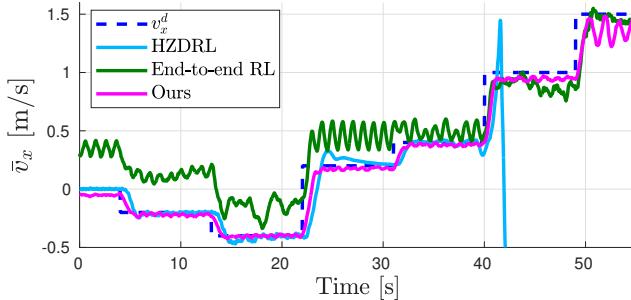


Fig. 8: Comparison of speed tracking performance with SOTA RL-based approaches.

In terms of sample efficiency, Fig. 9 shows our approach uses fewer data samples to successfully train a robust policy when compared with the end-to-end learning approach. Although this is not a surprising result as several authors have studied the effects of task space learning in the data sample efficiency [14], [26], to the best of our knowledge,

our method is the only task space approach that learns to walk from scratch and has been successfully implemented in 3D bipedal robots without the need of previously computed reference trajectories, e.g., the gait library used in [14].

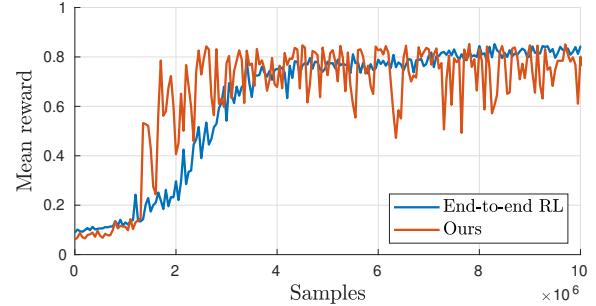


Fig. 9: Comparison of sample efficiency between traditional end-to-end RL and our proposed method. The higher variance in the reward is caused by the effect of the exploration noise in the task-space actions, which allows the policy to explore a more diverse set of behaviors during training.

D. Robustness on challenging terrain

We test our controller in terrains with slopes up to 20 degrees for Rabbit and Walker2D and 10 degrees for Digit. Fig. 10 shows a grid plot with the Root Mean Squared Error between \bar{v}_x and v_x^d for Digit with $v_x^d \in [-1.0, 1.0]\text{m/s}$, $v_y^d \in [-0.5, 0.5]\text{m/s}$, and $\alpha \in [-10, 10]$ degrees. We denote that during training, we only use $\alpha \in [0, 10]$. Yet, we test the policy in an extended range of slopes to highlight the robustness and interpolation capability of the policy to scenarios not seen during training. In this scenario, we define the robustness of the policy by its capability to keep track of v_x^d, v_y^d under challenging conditions. The results show that the learned policy keeps good tracking of v_x^d in almost all conditions, except for combinations where the robot is walking backward at high speed on very steep terrain.

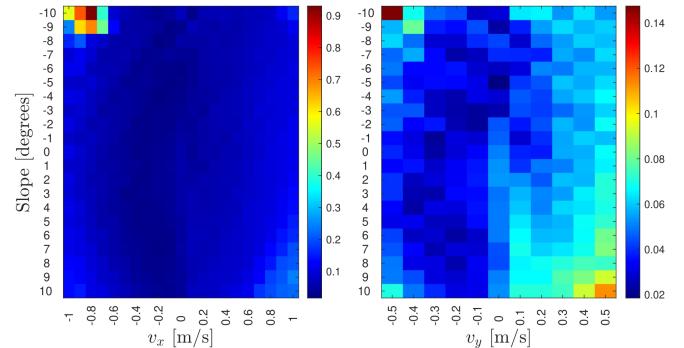


Fig. 10: Robustness of the learned policy under different conditions of terrain slope and target speed.

A higher error in the tracking of v_y^d is caused by the higher variance of the instantaneous velocity along the y axis, which is expected in bipedal locomotion as the robot's torso tends to oscillate more about the sagittal plane. We also notice there is a higher error when tracking positive speeds along the y axis. This effect may be caused by a bias created during training by the exposure of the policy to more samples with negative commands of v_y .

E. Robustness against disturbances.

The trained policy is also subjected to extensive tests against external disturbances $F_x \in [-80, 60]\text{N}$ applied in both forward and backward directions with duration $t \in [0.15, 3]\text{s}$. Fig. 11 shows the policy reacts effectively to all the applied disturbances without falling and maintaining the tracking of the desired walking speeds. In some scenarios, the policy uses interesting combinations of the task space outputs to realize effective strategies to reject the disturbances, e.g., bending forward/backward to absorb the impact. The results are similar for Walker2D and Digit. More details can be seen in the accompanying video submission.

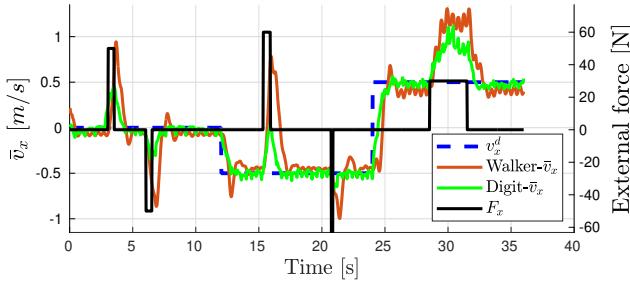


Fig. 11: Robustness test to external disturbances.

VI. CONCLUSION

In this work, we present a simple and effective learning-based hierarchical approach to realize robust locomotion controllers for bipedal robots. The design of the HL policy is inspired by the reduced-order state of the ALIP model, and a set of task space commands that include the step length, torso orientation, and height. This insightful choice of state and action spaces results in a compact policy that learns effective strategies for robust and dynamic locomotion. We show the HL planner is agnostic to the choice of LL controller, and its application is general to underactuated and fully actuated 2D and 3D robots. Finally, we show the learned policy tracks a wide range of speeds even under challenging conditions, such as external forces and slopes up to 20 degrees. Future work will focus on hardware experiments with the robot Digit and the extension of the proposed hierarchical framework to integrate different behaviors such as balancing, climbing stairs, and walking over stepping stones.

REFERENCES

- [1] A. Hereid, C. M. Hubicki, E. A. Cousineau, and A. D. Ames, “Dynamic humanoid locomotion: a scalable formulation for HZD gait optimization,” *IEEE Transactions on Robotics*, vol. 34, no. 2, pp. 370–387, Apr. 2018.
- [2] J. W. Grizzle, C. Chevallereau, R. W. Sinnet, and A. D. Ames, “Models, feedback control, and open problems of 3D bipedal robotic walking,” *Automatica*, vol. 50, no. 8, pp. 1955–1988, 2014.
- [3] Y. Gong and J. W. Grizzle, “Zero dynamics, pendulum models, and angular momentum in feedback control of bipedal locomotion,” *Journal of Dynamic Systems, Measurement, and Control*, vol. 144, no. 12, 10 2022.
- [4] G. Gibson, O. Dosunmu-Ogunbi, Y. Gong, and J. Grizzle, “Terrain-adaptive, alip-based bipedal locomotion controller via model predictive control and virtual constraints,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022.
- [5] X. B. Peng, G. Berseth, K. Yin, and M. Van De Panne, “Deeploco: Dynamic locomotion skills using hierarchical deep reinforcement learning,” *ACM Transactions on Graphics (TOG)*, vol. 36, 2017.
- [6] X. B. Peng, P. Abbeel, S. Levine, and M. van de Panne, “Deepmimic: Example-guided deep reinforcement learning of physics-based character skills,” *ACM Transactions on Graphics (TOG)*, vol. 37, no. 4, pp. 1–14, 2018.
- [7] S. Starke, I. Mason, and T. Komura, “Deepphase: Periodic autoencoders for learning motion phase manifolds,” *ACM Trans. Graph.*, vol. 41, no. 4, jul 2022.
- [8] Z. Xie, G. Berseth, P. Clary, J. Hurst, and M. van de Panne, “Feedback control for cassie with deep reinforcement learning,” in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1241–1246.
- [9] J. Siekmann, S. Valluri, J. Dao, L. Bermillo, H. Duan, A. Fern, and J. Hurst, “Learning memory-based control for human-scale bipedal locomotion,” *arXiv preprint arXiv:2006.02402*, 2020.
- [10] 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)*, 2021.
- [11] K. Green, Y. Godse, J. Dao, R. L. Hatton, A. Fern, and J. Hurst, “Learning spring mass locomotion: Guiding policies with a reduced-order model,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3926–3932, 2021.
- [12] G. A. Castillo, B. Weng, A. Hereid, Z. Wang, and W. Zhang, “Reinforcement learning meets hybrid zero dynamics: A case study for rabbit,” in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 284–290.
- [13] I. D. J. Rodriguez, N. Csomay-Shanklin, Y. Yue, and A. D. Ames, “Neural Gaits: Learning Bipedal Locomotion via Control Barrier Functions and Zero Dynamics Policies,” *arXiv*, 2022.
- [14] H. Duan, J. Dao, K. Green, T. Apgar, A. Fern, and J. Hurst, “Learning task space actions for bipedal locomotion,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [15] 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.
- [16] G. Castillo, B. Weng, W. Zhang, and A. Hereid, “Hybrid zero dynamics inspired feedback control policy design for 3d bipedal locomotion using reinforcement learning,” in *IEEE International Conference on Robotics and Automation (ICRA)*. Paris, France: IEEE, May 2020.
- [17] G. A. Castillo, B. Weng, W. Zhang, and A. Hereid, “Robust feedback motion policy design using reinforcement learning on a 3d digit bipedal robot,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 5136–5143.
- [18] 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*, Oct. 2012, pp. 5026–5033.
- [19] C. Chevallereau, G. Abba, Y. Aoustin, F. Plestan, E. R. Westervelt, C. C. De-Wit, and J. W. Grizzle, “RABBIT: a testbed for advanced control theory,” *IEEE Control Systems*, vol. 23, no. 5, Oct. 2003.
- [20] J. Reher, C. Kann, and A. D. Ames, “An Inverse Dynamics Approach to Control Lyapunov Functions,” *2020 American Control Conference (ACC)*, vol. 00, pp. 2444–2451, 2020.
- [21] A. Del Prete, F. Nori, G. Metta, and L. Natale, “Prioritized motion–force control of constrained fully-actuated robots: “task space inverse dynamics”,” *Robotics and Autonomous Systems*, vol. 63, pp. 150–157, 2015.
- [22] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *arXiv preprint arXiv:1707.06347*, 2017.
- [23] B. C. Bennett, S. D. Russell, P. Sheth, and M. F. Abel, “Angular momentum of walking at different speeds,” *Human Movement Science*, vol. 29, no. 1, pp. 114–124, Feb. 2010.
- [24] J. Siekmann, Y. Godse, A. Fern, and J. Hurst, “Sim-to-real learning of all common bipedal gaits via periodic reward composition,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 7309–7315.
- [25] H. Duan, A. Malik, J. Dao, A. Saxena, K. Green, J. Siekmann, A. Fern, and J. Hurst, “Sim-to-real learning of footstep-constrained bipedal dynamic walking,” in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 10428–10434.
- [26] G. Bellegarda and K. Byl, “Training in task space to speed up and guide reinforcement learning,” in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 2693–2699.