

Whole-body Humanoid Robot Locomotion with Human Reference

Qiang Zhang^{1,*}, Peter Cui^{2,*}, David Yan², Jingkai Sun¹, Yiqun Duan³, Gang Han², Wen Zhao²,
Weining Zhang², Yijie Guo², Arthur Zhang², Renjing Xu^{1,†}

Abstract—Recently, humanoid robots have made significant advances in their ability to perform challenging tasks due to the deployment of Reinforcement Learning (RL), however, the inherent complexity of humanoid robots, including the difficulty of designing complicated reward functions and training entire sophisticated systems, still poses a notable challenge. To conquer these challenges, after many iterations and in-depth investigations, we have meticulously developed a full-size humanoid robot, “Adam”, whose innovative structural design greatly improves the efficiency and effectiveness of the imitation learning process. In addition, we have developed a novel imitation learning framework based on an adversarial motion prior, which applies not only to Adam but also to humanoid robots in general. Using the framework, Adam can exhibit unprecedented human-like characteristics in locomotion tasks. Our experimental results demonstrate that the proposed framework enables Adam to achieve human-comparable performance in complex locomotion tasks, marking the first time that human locomotion data has been used for imitation learning in a full-size humanoid robot. For more video demonstrations, please visit our YouTube channel: <https://www.youtube.com/watch?v=7hK2ySYBa1I>

I. INTRODUCTION

In recent years, the field of humanoid robotics has received widespread attention, with numerous research institutes and companies releasing cutting-edge innovations and research results successively, signifying the rapid development and rise of the field. Boston Dynamics’ Atlas robot¹ has demonstrated parkour-level mobility; Tesla’s Optimus [1] and Figure’s humanoid robots² have learned from human data to perform complex desktop manipulation tasks; the bipedal robot Cassie [2] [3] [4] and its humanoid version Digit³, are powered by motors and successfully move across a variety of terrains; the renowned legged robot company Unitree released their humanoid robot product H1⁴; Apptroik⁵ has developed a humanoid robot named Apollo, powered entirely by push rod electric motors; OpenAI⁶, esteemed in the field of general-purpose AI, has acquired the 1X robot company⁷ and proposed a development plan for embodied intelligence. The above indicates that humanoid robotics is becoming one of the key directions for researchers and companies and that mastering the core technology of humanoid robotics is

¹The authors are with The Hong Kong University of Science and Technology (Guangzhou), China. * are equal contributors, † is the corresponding author {qzhang749, jsun444}@connect.hkust-gz.edu.cn, renjingxu@hkust-gz.edu.cn

²The authors are with PNDbotics. {peter.cui, david.yan, arthur.zhang}@pndbotics.com

³The author is with Human Centric AI Centre, Australia Artificial Intelligence Institute, University of Technology Sydney 2007 Ultimo Australia. yiqun.duan@student.uts.edu.au



Fig. 1: The top image displays the humanoid robot Adam walking on unseen terrain, and the bottom image shows Adam moving from a standing position to running.

crucial to bridging the gap between digital general-purpose AI and tangible hardware.

Traditional robot control algorithms, which typically rely on precise mathematical models and predefined motion planning, have proven to be very effective in the past for locomotion tasks of quadrupedal, bipedal, and humanoid robots. Notably, Boston Dynamics’ Atlas and Spot robots have demonstrated the effectiveness of these methods by using Model Predictive Control (MPC) algorithms [5] to present impressive extreme mobility in a variety of demonstrations. However, these algorithms often rely on accurate modeling of the environment, which can introduce significant challenges in terms of robustness and generalisability, especially in unknown or dynamically changing environments, where the performance of traditional control algorithms can be significantly degraded, limiting their usefulness in a wider range of application scenarios. This reliance on accurate modeling, moreover, requires a high level of expertise to build and maintain these models, increasing the complexity of development and debugging.

Traditional robot control algorithms are notably limited in their adaptability, flexibility, and user-friendliness even though they exhibit excellent performance in specific environments, motivating researchers to explore alternative approaches to overcome these obstacles to design smarter and more adaptable robot control strategies. Among them, deep neural network-based reinforcement learning algorithms have

achieved promising results in the control of legged robots [6]. By interacting with the environment, Deep reinforcement learning algorithms are able to autonomously discover effective strategies to perform complex tasks and can be potentially extended to unknown or dynamically changing environments, which provides robots with unprecedented adaptability and flexibility.

Reinforcement learning algorithms have made significant progress in a variety of legged robots, but applications in the field of humanoid robotics still lack sufficient exploration. This can be mainly attributed to the following factors: Firstly, most humanoid robots are expensive and difficult to maintain, which is a considerable obstacle for research institutes with limited funds; secondly, the interpretability problem of deep neural networks and the Sim2Real gap in the training process of deep reinforcement learning make it difficult to transfer the models to practical applications; Lastly, the complexity of humanoid robots greatly exceeds that of other legged robots, which makes the design of reward functions and training strategies during their training more challenging.

To address these challenges, after repeated design and in-depth exploration, we have introduced the motor-joint-driven humanoid robot “Adam”, which has a significant cost advantage over traditional hydraulically-driven robots, and its modular design facilitates repairs during experiments and further reduces maintenance costs. Moreover, our high-performance actuators ensure the robot’s exceptional mobility, granting it a range of motion in its limbs close to that of a human. For the difficulty of setting up complicated reward functions in reinforcement learning, we adopt an innovative strategy of using human motion data to guide the learning process. Combined with our imitation learning training framework, Adam’s performance on the locomotion task in the experiment is impressive.

In summary, we introduce a brand-new humanoid robot, Adam, and provide a new methodology and experimental validation for the learning, adaptation, and optimization of humanoid robots, paving a new way for research and development in humanoid robotics. The contributions of this paper can be summarized in the following three points:

- 1) We have developed and detailed an innovative biomimetic humanoid robot, Adam, whose limbs not only have a range of motion close to that of humans but also offer sufficient advantages in terms of low cost and ease of maintenance.
- 2) We designed and validated a new whole-body imitation learning framework for humanoid robots, which effectively solves the problem of complex reward function settings encountered in the reinforcement learning training of humanoid robots, greatly reduces the Sim2Real gap and improves the learning ability and adaptability of humanoid robots.
- 3) To address the Sim2Real challenges in complex humanoid robot reinforcement learning control algorithms, we incorporated numerous cross-validation and feedback adjustment steps into our framework. We not only demonstrated the robot’s highly human-like

performance in executing complex motion tasks but also provided a new perspective and data support for the future motion learning and optimization of humanoid robots.

II. RELATED WORK

A. Legged robots locomotion

Legged robot locomotion has seen significant advancement as it serves as the fundamental basis for achieving humanoid robotics. The incorporation [7], [8] of Reinforcement Learning (RL) has played a crucial role in this field. The robot Cassie realizes a wide range of walking and running patterns utilizing periodic-parametrized reward functions [9], and even set a Guinness World Record for the fastest 100m dash using pre-optimized gaits [10]. Jeon et al. [11] underscores the effectiveness of potential-based reward shaping in expediting learning and enhancing the robustness of legged locomotion. Similarly, Shi et al. [12] have broadened the capabilities of legged robots by introducing an assistive force curriculum that facilitates agile motion learning in settings without explicit references. The HRP-5P humanoid robot has demonstrated exceptional bipedal walking by leveraging feedback from actuator currents [13], while Kim et al. [14] have proposed a torque-based approach to bridge the gap between simulated training and real-world application effectively. In an innovative endeavor, DeepMind’s research [15] enabled a miniature humanoid robot to master complex soccer skills through a unique teacher-student distillation and self-play methodology. Additionally, utilizing attention-based transformers in the Digit humanoid robot has facilitated more adaptable and versatile locomotion patterns [16]. More recently, Tang et al. [17] employed adversarial critic component and specially designed Wasserstein distances to migrate locomotion from human reference. However, Adam has more flexibility and is more human-like overall, which leads to better performance in real robot experiments.

B. Learning from Human Reference

Humans, with their advanced intelligence and versatile locomotion capabilities, exhibit complex motion patterns that embody rich information. Leveraging insights from human behavior through learning from human references can greatly enhance the adaptability of robots. Traditional approaches to behavior cloning relied on manual programming, which proved to be time-consuming and inflexible [18], [19]. Furthermore, defining the intricate and versatile locomotion skills of humanoid robots manually posed significant challenges [20], [21].

In recent years, imitation learning (IL) strategies have gained prominence, involving the tracking of either reference joint trajectories or extracted gait features [22]–[25]. However, these tracking techniques often operate on individual motion clips, resulting in discontinuities when transitioning between different locomotion patterns. To address this limitation, Generative Adversarial Imitation Learning (GAIL) was introduced by Peng et al., who proposed two innovative methods: AMP and Successor ASE [26]–[28].

These approaches enable physics-based avatars to perform objective tasks while implicitly imitating diverse motion styles from extensive unstructured datasets. Variants of AMP have been successfully applied to learn agile quadrupedal locomotion and terrain-adaptive skills [29]–[33]. Additionally, Tang et al. introduced a Wasserstein adversarial imitation system with soft boundary constraints, further enhancing the capabilities of the AMP method [17]. To facilitate the transfer of reference motion to robots, many works have introduced re-targeting techniques [17], [34], [35] that consider primitive skeleton and geometry consistency, enabling accurate dynamic modeling and complex balance controllers. However, AMP only makes robots learn kinematic motion relationships, which lacks physics constraints and causes a huge sim-to-real gap.

III. PRELIMINARY

A. The Structure of Humanoid Robot Adam

In this paper, we conducted experiments using the lite version of Adam. Adam(Lite) is equipped with 25 QDD(quasi-direct drive), force-controlled PND actuators throughout its body, standing at 1.6 meters tall and weighing 60 kilograms. Its legs are fitted with four QDD high-sensitivity, highly back-drivable actuators with up to 340N·m maximum torque. The arms have five degrees of freedom, and the waist has three. This fully modular, highly reusable design of flexible actuators, along with a highly biomimetic torso configuration, provides Adam with exceptional mobility and adaptability. The entire body utilizes a full-stack in-house design, real-time communication network PDN (PND Network), and PND actuators. The motion control computer is the 12th generation Intel i7 processor (Intel NUC) and PND RCU(Robot control unit). The PND RCU integrates all actuators, battery management system (BMS), power management, and is equipped with a 16-port Gigabit Ethernet programmable switch with network management capabilities, forming the robot’s sensory and control communication hub. This configuration enables Adam to perform large-scale parallel dynamic simulations and neural network training, achieving diversified full-body motion control that is truly applicable to real service scenarios, adapting to the complex and variable human social environment. Dexterous hands and vision modules can be optionally equipped. Since we focused on blind locomotion tasks in this paper, these parts were not included. The schematic diagram of Adam(Lite)’s entire body structure is shown in Figure. 2, with information on its joint structure and more detailed in TABLE. I.

TABLE I: Specifications of Humanoid Robot Adam.

| Humanoid Robot | Adam (Lite) |
|-------------------------------|-----------------------------|
| Height | 1.6 M |
| Weight | 60 Kg |
| Full-body degrees of freedom | 25 |
| Single-leg degrees of freedom | Hip x3 + Knee x1 + Ankle x2 |
| Single-arm degrees of freedom | Shoulder x3 + Elbow x2 |
| Waist degrees of freedom | 3 |

B. Motion Capture and Re-Targeting

In our research, we explored a variety of human motion data sources to enrich our training set, ensuring that our




Fig. 2: Schematic representation of the humanoid robot Adam(Lite), showing both the front view and the angled side view. Particularly, the design of its hip joint distinctly emulates the human skeleton, demonstrating a remarkable level of anthropomorphism throughout its structure.

model could learn diverse human motion characteristics. Initially, we utilized two public motion capture (mocap) databases: the SFU mocap dataset ⁸ and the CMU mocap dataset ⁹. Both datasets contain multiple motion capture sequences, covering a wide range of human activities, including everyday actions, sports movements, dance, and combat actions. By integrating these two databases, we provided Adam with a diverse and high-quality human whole-body motion dataset, which is crucial for training the robot to understand and mimic human motion patterns.

Beyond public datasets, we also used high-precision motion capture equipment for custom motion recording. This approach allowed us to capture specific motion data, especially those special actions or sequences designed for particular experimental needs that are difficult to find in public databases. These custom motion data not only added diversity to our dataset but also enabled us to fine-tune and optimize our model more precisely, adapting it to specific motion tasks and challenges. We did not approach re-targeting as an optimization problem, as was done in approaches like humanmimic [17]. Our objective was to secure high-quality motion data tailored for Adam, which led us to manually calibrate and convert each dataset ourselves.

C. Reinforcement Learning on Legged Robots

The application of RL to legged robots is modeled as a Partially Observable Markov Decision Process (POMDP), defined by the tuple $(\mathcal{S}, \mathcal{O}, \mathcal{A}, \mathcal{R}, p, \gamma)$ where \mathcal{S} defines the state space, including all states of the environment. \mathcal{O} denotes the partial observation space, representing a subset of the state space \mathcal{S} that covers only the observable aspects of the environment by the agent. Concurrently, \mathcal{A} specifies the action space, indicating all actions the agent is capable of executing. The reward function, denoted by \mathcal{R} , assigns a scalar reward to each state-action pair. The transition probabilities, represented by p , specify the likelihood of transition from the current state to a new state following the chosen action. The discount factor $\gamma \in [0, 1]$, is applied to ensure that future rewards are appropriately weighted.

Within this framework, at each timestep t , the agent observes $o_t \in \mathcal{O}$ from the environment. Based on this ob-

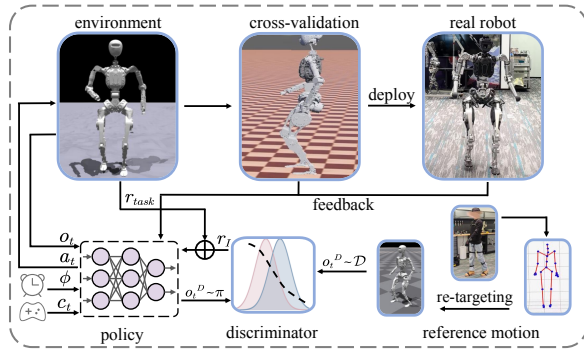


Fig. 3: Adversarial Motion Priors Imitation Training Framework of Humanoid Robot

ervation, the agent outputs an action $a_t \in \mathcal{A}$ sampled from a policy $\pi(a_t|o_t)$. Subsequently, the environment transitions to a new state s_{t+1} , governed by the probability distribution $s_{t+1} \sim p(s_{t+1}|s_t, a_t)$, and the agent is accordingly awarded a reward $r_t = \mathcal{R}(s_t, a_t)$. This objective function is formalized as to maximize the received reward:

$$\arg \max_{\theta} \mathbb{E}_{(s_t, a_t) \sim p_{\theta}(s_t, a_t)} \left[\sum_{t=0}^{T-1} \gamma^t r_t \right] \quad (1)$$

where T denotes the time horizon of POMDP.

IV. METHOD

A. Learning with Human Reference by Adversarial Motion Priors

Our humanoid imitation framework is built upon the AMP. In our framework, the discriminator D outputs the similarity between the state transitions sampled from the agent and it sampled from the reference demonstrations \mathcal{D} . It is crucial to select the discriminator observation $o_t^D \in \mathbb{R}^{58}$ fed into the discriminator to ensure the robots with similar state transitions can perform similar locomotion styles. The discriminator observation contains the velocity of each actuated joint and the positions of two hands and two feet of the humanoid. At each time step, we randomly sample the state transitions from the demonstrations and feed them into the discriminator to get the expert prediction loss for enabling the discriminator to distinguish them,

$$\mathcal{L}_{expert} = \mathbb{E}_{(o_t^D, o_{t+1}^D) \sim \mathcal{D}} [(D(o_t^D, o_{t+1}^D) - 1)^2] \quad (2)$$

As do the same with state transitions sampled from the policy,

$$\mathcal{L}_{policy} = \mathbb{E}_{(o_t^D, o_{t+1}^D) \sim \pi} [(D(o_t^D, o_{t+1}^D) + 1)^2] \quad (3)$$

We follow to [27] to formulate the penalty for gradients on samples from the reference to stabilize training,

$$\mathcal{L}_{GP} = \mathbb{E}_{(o_t^D, o_{t+1}^D) \sim \mathcal{D}} [\|\nabla \mathcal{D}(o_t^D, o_{t+1}^D)\|^2] \quad (4)$$

Finally, we formulate the total AMP loss as,

$$\mathcal{L}_{AMP} = \frac{1}{2} \mathcal{L}_{expert} + \frac{1}{2} \mathcal{L}_{policy} + \lambda_{GP} \mathcal{L}_{GP} \quad (5)$$

The AMP loss function directs the discriminator to rate samples, giving scores close to +1 for genuine reference motions and nearing -1 for those generated by the policy.




Fig. 4: Visualization of periodic rewards based on the von Mises distribution

The objective of the policy is to create motions convincing enough to lead the discriminator into assigning higher scores, demonstrating its ability to closely mimic reference motions. Subsequently, the formulation of the imitation reward for the policy's training is denoted as,

$$r_I = \max[0, 1 - \frac{1}{4}(D(o_t^D, o_{t+1}^D) - 1)^2] \quad (6)$$

where o_t^D, o_{t+1}^D are sampled from the policy.

B. End-to-End Reinforcement Learning on Humanoid Robot

Concurrently, the movement direction in reference motion is typically limited to the local coordinate. To facilitate control under the world coordinate system, generate more natural gaits, and enable a more effective transition from simulation to reality on challenging terrains, we introduced coordinated task rewards. The task rewards are composed of three parts, the command reward, the periodic reward, and the regularization reward. The command reward forces the robot to move alone in the command directions, which is formulated as,

$$r_{com} = \sum \lambda_i \exp(-\omega_i (|\mathbf{v}_{des}^i - \mathbf{v}_t^i|)) \quad i \in (x, y, yaw) \quad (7)$$

where λ_i and ω_i are the weight of each direction of command rewards, \mathbf{v}_{des} is the desire velocity vector along the specific direction and \mathbf{v}_t is the velocity vector at t .

To facilitate the attainment of the desired gait performance, we introduce periodic rewards aligned with imitation rewards. This method naturally promotes the maintenance of a stable gait in the robot. However, if a variable gait is desired, it would be advisable to omit this reward function. We follow [9] to formulate the periodic reward by the swing phases, characterized by the foot moving through the air, and the stance phases, where the foot is firmly planted on the ground. Each periodic reward item is composed of a coefficient α_i , a phase indicator $I_i(\phi)$, a valued phase reward function $V_i(s_t)$, ϕ is the cycle time, i denotes the phase is stance phase or swing phase. The swing and stance phases are sequentially arranged and collectively span the entire cycle duration by establishing a ratio $\rho \in (0, 1)$. This configuration ensures that the swing phase occupies a duration equivalent to ρ , followed immediately by the stance phase which extends for a duration of $1 - \rho$. The single-foot reward is written as follows,

$$\begin{aligned} r_{per} &= \sum \alpha_i \mathbb{E}[I_i(\phi)] V_i(s_t) \\ V_{stance}(s_t) &= \exp(-10F_f^2) \\ V_{swing}(s_t) &= \exp(-200v_f^2) \end{aligned} \quad (8)$$

where F_f is the norm force of each foot, v_f is the speed of each foot. For modeling the phase indicator $I_i(\phi)$, we follow the [36] to use the mathematical expectation of Von Mises distribution. The visualization of the phase indicator is shown in Figure 4. And we formulate as,

$$\begin{aligned} Q_1 &= I_{stance}(\phi + \theta_{left}) \\ Q_2 &= I_{stance}(\phi + \theta_{right}) \end{aligned} \quad (9)$$

where the θ_{left} , θ_{right} is the offsets of left and right leg in cycle time. Expect the normal periodic reward, we calculate rewards for foot speed, height difference, and the symmetric during the swing phase for more natural gait style. The foot speed tracking reward is formulated as,

$$\begin{aligned} q^i &= \text{clip}\left(\frac{\phi}{\rho} - 0.5, 0, 1\right) \\ r(s_t) &= \begin{cases} 16(q^i v_f^i)^2 & 0 \leq q_i \leq 0.6, \\ 0 & q_i > 0.6. \end{cases} \end{aligned} \quad (10)$$

where $i \in (\text{left}, \text{right})$, ‘‘clip’’ is the function to constrain the values of variables within the range of 0 to 1. The foot speed tracking reward encourages the robots to perform higher foot speed during the swing phase. The height difference reward is,

$$\begin{aligned} q^i &= \frac{\phi}{\rho} \\ \delta h &= h_f^i - h_f^{-i} - 0.02 \\ r(s_t) &= \begin{cases} 2\exp(-25(|\delta h|)) & 0 \leq q_i \leq 0.3, \\ 0 & q_i > 0.3. \end{cases} \end{aligned} \quad (11)$$

where h_f^i is the height i foot, the h_f^{-i} denotes the height of the other side foot. The purpose of this function is to calculate reward based on the height difference of the foot only during certain early stages of the gait cycle. The symmetric reward is written as follows,

$$\begin{aligned} \mathbf{d}_t &= \mathbf{p}_t^{left} - \mathbf{p}_t^{right} \\ tf &= (\mathbb{E}[I_{left}(\phi)] > 0.5) \wedge (\mathbb{E}[I_{right}(\phi)] > 0.5) \\ \delta \mathbf{f}_t &= tf \cdot \mathbf{d}_t + \neg tf \cdot \delta \mathbf{f}_{t-1} \\ \delta \mathbf{l}_t &= \neg tf \cdot \delta \mathbf{f}_t + tf \cdot \mathbf{d}_t \\ r(s_t) &= 3.3tf \exp(-10\|\mathbf{d}_t + \delta \mathbf{l}_t\|_1) \end{aligned} \quad (12)$$

where \mathbf{p}_t^i is the 3-D position of foot end-effector, \neg is negation symbol. This function calculates a reward based on the symmetry of the gait swing, where the gait phases of the left and right feet are taken into consideration, along with the distance between the feet.

To enhance the robustness of the sim-to-real transfer, we incorporated regularization rewards into the comprehensive reward structure. These rewards enforce motion constraints, emphasizing both smoothness and safety. The detailed formulation of each reward is shown in Table II. The DoF limits reward and keeps the movement within the robot’s physical capabilities, ensuring that it doesn’t attempt movements that could damage its mechanisms or be outside its operational scope. \mathbf{b} denotes the vector angle of DoF. \mathbf{b}_{lower} and \mathbf{b}_{upper} are the upper and lower bounds of the joint limit. The DoF velocity reward forces the system to maintain optimal speeds. Similarly, the DoF acceleration reward encourages smooth

TABLE II: Regularization Rewards

| Item | Detail |
|--------------------------|---|
| Action differential | $\exp(-0.05\ \mathbf{a}_t - \mathbf{a}_{t-1}\ _2)$ |
| DoF limits | $\exp(-2.0(\min(0, \mathbf{b} - \mathbf{b}_{upper}) - \max(0, \mathbf{b} - \mathbf{b}_{lower})))$ |
| DoF velocity | $\exp(-1 \times 10^{-4}\ \dot{\mathbf{b}}\ _2^2)$ |
| DoF acceleration | $\exp(-1 \times 10^{-7}\ \ddot{\mathbf{b}}\ _2^2)$ |
| Arm DoF penalty | $\exp(\ \mathbf{b}_{arm}\ _1)$ |
| Orientation differential | $\exp(-300(\text{roll}^2 + \text{pitch}^2))$ |
| Torso yaw | $\ \theta_{torso\ yaw}\ _1$ |
| Torques | $\exp(-5 \times 10^{-4}\ \tau_t\ _2)$ |

TABLE III: Domain Randomization Range

| Randomization Item | Range | Unit |
|------------------------|-------------------|------|
| Mass | [-0.05,0.05] | Kg |
| Center of mass alone x | [-0.05,0.05] | m |
| Center of mass alone z | [-0.05,0.05] | m |
| Motor strength | [0.7,1.4]×default | N·m |
| Impulse | [0,0.8] | m/s |
| External force | [-500,500] | N |
| Linear velocity | [0.8,1.2]×default | m/s |

increases and decreases in speed, contributing to the overall stability. For robotic arms, the arm DoF penalty is applied to discourage large positions or movements. The torso yaw reward and orientation differential rewards are specific to the rotational movement around a robot’s vertical axis. This control is vital for balance and orientation, especially in humanoid robots.

V. EXPERIMENTS

A. Training and Implementation Details

1) *Training Setup*: We trained our policy via a model-free reinforcement learning algorithm known as Proximal Policy Optimization (PPO) [37] on 4096 Isaac Gym simulation environments parallel. In humanoid locomotion scenarios, actions are represented by $\mathbf{a}_t \in \mathbb{R}^{25}$ a 25-dimensional vector specifying the desired positional adjustments for each actuated joint as dictated by the Proportional-Derivative (PD) controller. Observations, denoted as $\mathbf{o}_t \in \mathbb{R}^{91}$, encompass the humanoid’s current linear and angular velocities, the average velocities in the (x, y, yaw) directions, and the orientation of the gravity vector within the robot’s base frame, all measured by an inertial measurement unit (IMU). The position and velocity of each joint are captured by the actuators’ encoders. A command, $c_t = (v_{des}^x, v_{des}^y, \omega_{des}^{yaw})$, is introduced to specify the desired velocities along the x-axis, y-axis, and yaw within the robot’s base frame. Periodic observations include the sine and cosine of the cycle time and the swing phase ratio ρ , to dynamically adjust the gait. Additionally, the root’s height is incorporated into the observations to manage the humanoid robot’s posture. For enhanced performance, the action executed in the previous timestep is also considered. Our framework ensures a smooth training process for Adam, as illustrated in Figure 6.

2) *Domain Randomization*: This randomization strategy aims to tackle three critical uncertainties: modeling discrepancies in the robot and its environment, sensor noise, and unobservable disturbances and impulses. To avoid the sim-to-real gap in force applying in Isaac Gym, we use the sudden velocity on the root to replace the force impulse. These dynamic properties are represented in Table III, with

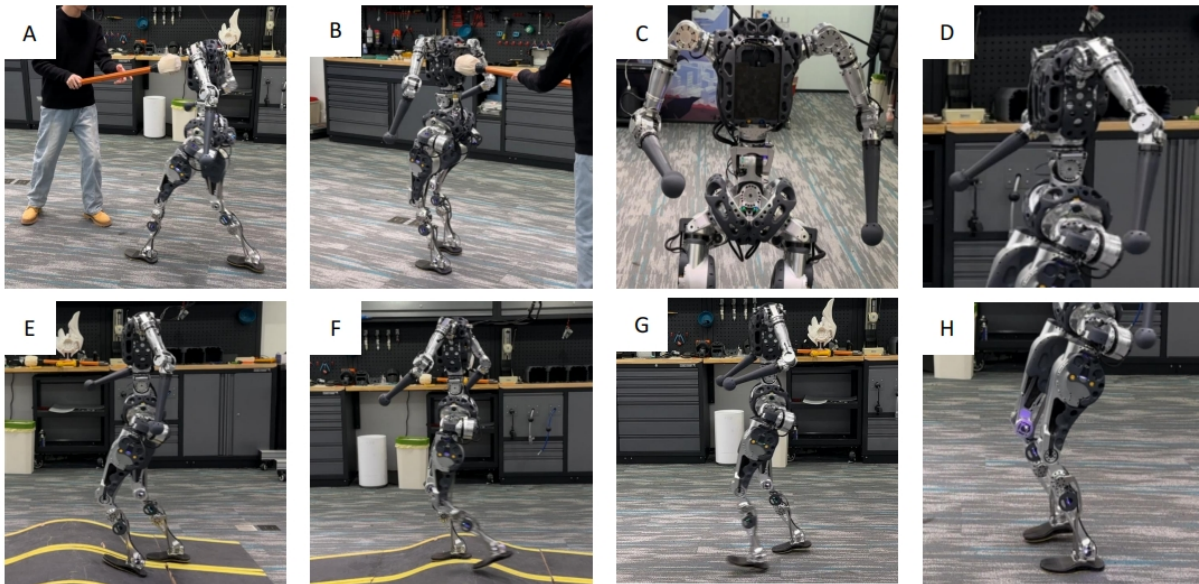


Fig. 5: **Real Robot Experiments.** We tested our method on Adam. (A)(B) demonstrate the robot’s robust locomotion performance under external disturbances. **Especially, it demonstrated the characteristic of “straight-knee”** (C)(D) show the human-like swinging of the upper limbs. (E)(F) display the robustness and human-likeness on unknown and complex terrains. (G)(H) present for the first time the humanoid robot’s “heel-to-toe” running and walking gaits.




Fig. 6: Curve of normalized return and episode length of policy training




Fig. 7: Cross-validation of Webots and Isaac gym simulator in forward locomotion scenario

values adjusted across the specified ranges. The mass of individual robot bodies (*e.g.*, torso), is varied to simulate different weight distributions. The center of mass is randomly shifted in specific directions. These techniques help the policy to adapt to the uncertainties in robot components. Motor strength influences the torque output to deal with the uncertainties in motors. Furthermore, impulses are applied by adding velocities in the x - y plane on the robot root, replicating the effect of sudden external shocks. External forces are also directly applied to the body of the robot to simulate unobservable environmental forces. Additionally, noise is injected into the linear velocity to simulate speed estimation errors. A series of tests were performed under a

variety of conditions and in response to external disturbances, with the outcomes illustrated in Figure. 5. These results shows our method provided a new perspective of future motion learning for humanoid robots

B. Cross-validation and Feedback Fine-tuning

We carried out comprehensive cross-validation experiments using both Webots¹⁰ and Isaac Gym as simulation platforms to evaluate the effectiveness of our models in environments that closely mimic real-world conditions. Notably, Webots offers advanced physical simulation capabilities, enabling us to validate our policy across diverse platforms while fine-tuning hyperparameters for optimal performance. Which vividly demonstrated in Figure 7, where the humanoid robot’s gait appears synchronized and uniform across both simulators, underscoring the model’s reliability. Additionally, we conducted real-world tests to further validate our model’s practical applicability. The seamless transition of the humanoid robot from a standing to a running state in the real world is captured in Figure 1, highlighting the model’s effectiveness and the potential for real-world deployment.

VI. CONCLUSION

In this paper, we implement a whole-body imitation learning framework on the humanoid robot Adam, whose performance in complex gaits tasks is comparable to that of humans, and demonstrate human-like features such as “heel-to-toe” transitions and straight-knee movements for the first time in humanoid robots. Our experiments have demonstrated the practicality and efficiency of our framework, highlighting its great potential for further research in humanoid robotics. Moreover, our framework supports the integration of perceptual modules, and in the future, we plan to incorporate more sensors to allow Adam to imitate human motions under integrated perceptual conditions.

REFERENCES

- [1] Wikipedia, “Optimus (robot),” [https://en.wikipedia.org/wiki/Optimus_\(robot\)](https://en.wikipedia.org/wiki/Optimus_(robot)), 2024, accessed: 2024-02-28.
- [2] Y. Gong, R. Hartley, X. Da, A. Hereid, O. Harib, J.-K. Huang, and J. Grizzle, “Feedback control of a cassie bipedal robot: Walking, standing, and riding a segway,” in *2019 American Control Conference (ACC)*. IEEE, 2019, pp. 4559–4566.
- [3] Z. Li, X. B. Peng, P. Abbeel, S. Levine, G. Berseth, and K. Sreenath, “Reinforcement learning for versatile, dynamic, and robust bipedal locomotion control,” *arXiv preprint arXiv:2401.16889*, 2024.
- [4] 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 *IEEE International Conference on Robotics and Automation (ICRA)*, Xi’an, China, June 2021.
- [5] C. E. Garcia, D. M. Prett, and M. Morari, “Model predictive control: Theory and practice—a survey,” *Automatica*, vol. 25, no. 3, pp. 335–348, 1989.
- [6] N. Rudin, D. Hoeller, P. Reist, and M. Hutter, “Learning to walk in minutes using massively parallel deep reinforcement learning,” in *Conference on Robot Learning*. PMLR, 2022, pp. 91–100.
- [7] Z. Li, X. B. Peng, P. Abbeel, S. Levine, G. Berseth, and K. Sreenath, “Robust and versatile bipedal jumping control through reinforcement learning.”
- [8] J. Siekmann, K. Green, J. Warila, A. Fern, and J. Hurst, “Blind bipedal stair traversal via sim-to-real reinforcement learning,” *arXiv preprint arXiv:2105.08328*, 2021.
- [9] J. Siekmann, Y. Godse, A. Fern, and J. Hurst, “Sim-to-real learning of all common bipedal gaits via periodic reward composition,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 7309–7315.
- [10] D. Crowley, J. Dao, H. Duan, K. Green, J. Hurst, and A. Fern, “Optimizing bipedal locomotion for the 100m dash with comparison to human running,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 12 205–12 211.
- [11] S. H. Jeon, S. Heim, C. Khazoom, and S. Kim, “Benchmarking potential based rewards for learning humanoid locomotion,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 9204–9210.
- [12] F. Shi, Y. Kojio, T. Makabe, T. Anzai, K. Kojima, K. Okada, and M. Inaba, “Reference-free learning bipedal motor skills via assistive force curricula,” in *The International Symposium of Robotics Research*. Springer, 2022, pp. 304–320.
- [13] R. P. Singh, Z. Xie, P. Gergondet, and F. Kanehiro, “Learning bipedal walking for humanoids with current feedback,” *IEEE Access*, 2023.
- [14] 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,” *arXiv preprint arXiv:2304.09434*, 2023.
- [15] T. Haarnoja, B. Moran, G. Lever, S. H. Huang, D. Tirumala, M. Wulfmeier, J. Humplik, S. Tunyasuvunakool, N. Y. Siegel, R. Hafner, *et al.*, “Learning agile soccer skills for a bipedal robot with deep reinforcement learning,” *arXiv preprint arXiv:2304.13653*, 2023.
- [16] I. Radosavovic, T. Xiao, B. Zhang, T. Darrell, J. Malik, and K. Sreenath, “Learning humanoid locomotion with transformers,” *arXiv preprint arXiv:2303.03381*, 2023.
- [17] A. Tang, T. Hiraoka, N. Hiraoka, F. Shi, K. Kawaharazuka, K. Kojima, K. Okada, and M. Inaba, “Humanmimic: Learning natural locomotion and transitions for humanoid robot via wasserstein adversarial imitation,” *arXiv preprint arXiv:2309.14225*, 2023.
- [18] T. Osa, J. Pajarinen, G. Neumann, J. A. Bagnell, P. Abbeel, J. Peters, *et al.*, “An algorithmic perspective on imitation learning,” *Foundations and Trends® in Robotics*, vol. 7, no. 1-2, pp. 1–179, 2018.
- [19] H. Ravichandar, A. S. Polydoros, S. Chernova, and A. Billard, “Recent advances in robot learning from demonstration,” *Annual review of control, robotics, and autonomous systems*, vol. 3, pp. 297–330, 2020.
- [20] S. Bohez, S. Tunyasuvunakool, P. Brakel, F. Sadeghi, L. Hasenclever, Y. Tassa, E. Parisotto, J. Humplik, T. Haarnoja, R. Hafner, *et al.*, “Imitate and repurpose: Learning reusable robot movement skills from human and animal behaviors,” *arXiv preprint arXiv:2203.17138*, 2022.
- [21] L. Han, Q. Zhu, J. Sheng, C. Zhang, T. Li, Y. Zhang, H. Zhang, Y. Liu, C. Zhou, R. Zhao, *et al.*, “Lifelike agility and play on quadrupedal robots using reinforcement learning and generative pre-trained models,” *arXiv preprint arXiv:2308.15143*, 2023.
- [22] S. Schaal, “Is imitation learning the route to humanoid robots?” *Trends in cognitive sciences*, vol. 3, no. 6, pp. 233–242, 1999.
- [23] J. Van Den Berg, S. Miller, D. Duckworth, H. Hu, A. Wan, X.-Y. Fu, K. Goldberg, and P. Abbeel, “Superhuman performance of surgical tasks by robots using iterative learning from human-guided demonstrations,” in *2010 IEEE International Conference on Robotics and Automation*. IEEE, 2010, pp. 2074–2081.
- [24] P. M. Kebria, A. Khosravi, S. M. Salaken, and S. Nahavandi, “Deep imitation learning for autonomous vehicles based on convolutional neural networks,” *IEEE/CAA Journal of Automatica Sinica*, vol. 7, no. 1, pp. 82–95, 2020.
- [25] L. Le Mero, D. Yi, M. Dianati, and A. Mouzakitis, “A survey on imitation learning techniques for end-to-end autonomous vehicles,” *IEEE Transactions on Intelligent Transportation Systems*, 2022.
- [26] J. Ho and S. Ermon, “Generative adversarial imitation learning,” *Advances in neural information processing systems*, vol. 29, 2016.
- [27] X. B. Peng, Z. Ma, P. Abbeel, S. Levine, and A. Kanazawa, “Amp: Adversarial motion priors for stylized physics-based character control,” *ACM Transactions on Graphics (ToG)*, vol. 40, no. 4, pp. 1–20, 2021.
- [28] X. B. Peng, Y. Guo, L. Halper, S. Levine, and S. Fidler, “Ase: Large-scale reusable adversarial skill embeddings for physically simulated characters,” *ACM Transactions On Graphics (TOG)*, vol. 41, no. 4, pp. 1–17, 2022.
- [29] A. Escontrela, X. B. Peng, W. Yu, T. Zhang, A. Iscen, K. Goldberg, and P. Abbeel, “Adversarial motion priors make good substitutes for complex reward functions,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 25–32.
- [30] E. Vollenweider, M. Bjelonic, V. Klemm, N. Rudin, J. Lee, and M. Hutter, “Advanced skills through multiple adversarial motion priors in reinforcement learning,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 5120–5126.
- [31] C. Li, M. Vlastelica, S. Blaes, J. Frey, F. Grimmering, and G. Martius, “Learning agile skills via adversarial imitation of rough partial demonstrations,” in *Conference on Robot Learning*. PMLR, 2023, pp. 342–352.
- [32] J. Wu, G. Xin, C. Qi, and Y. Xue, “Learning robust and agile legged locomotion using adversarial motion priors,” *IEEE Robotics and Automation Letters*, 2023.
- [33] Y. Wang, Z. Jiang, and J. Chen, “Amp in the wild: Learning robust, agile, natural legged locomotion skills,” *arXiv preprint arXiv:2304.10888*, 2023.
- [34] K. Ayusawa and E. Yoshida, “Motion retargeting for humanoid robots based on simultaneous morphing parameter identification and motion optimization,” *IEEE Transactions on Robotics*, vol. 33, no. 6, pp. 1343–1357, 2017.
- [35] R. Grandia, F. Farshidian, E. Knoop, C. Schumacher, M. Hutter, and M. Bächer, “Doc: Differentiable optimal control for retargeting motions onto legged robots,” *ACM Transactions on Graphics (TOG)*, vol. 42, no. 4, pp. 1–14, 2023.
- [36] J. H. Park, “Impedance control for biped robot locomotion,” *IEEE Transactions on Robotics and Automation*, vol. 17, no. 6, pp. 870–882, 2001.
- [37] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *arXiv preprint arXiv:1707.06347*, 2017.

NOTES

¹<https://bostondynamics.com/atlas/>

²<https://www.figure.ai/>

³<https://agilityrobotics.com/robots>

⁴<https://www.unitree.com/h1/>

⁵<https://apptronik.com/>

⁶<https://openai.com/>

⁷<https://www.1x.tech/>

⁸<https://mocap.cs.sfu.ca/>

⁹<http://mocap.cs.cmu.edu/>

¹⁰ <https://cyberbotics.com/>