

ECO: Energy-Constrained Optimization with Reinforcement Learning for Humanoid Walking

Weidong Huang, Jingwen Zhang, *Member, IEEE*, Jiongye Li, Shibowen Zhang, Jiayang Wu, Jiayi Wang, Hangxin Liu, *Member, IEEE*, Yaodong Yang, *Member, IEEE*, Yao Su, *Member, IEEE*

Abstract—Achieving stable and energy-efficient locomotion is essential for humanoid robots to operate continuously in real-world applications. Existing model predictive control (MPC) and reinforcement learning (RL) approaches often rely on energy-related metrics embedded within a multi-objective optimization framework, which require extensive hyperparameter tuning and often result in suboptimal policies. To address these challenges, we propose ECO (Energy-Constrained Optimization), a constrained RL framework that separates energy-related metrics from rewards, reformulating them as explicit inequality constraints. This method provides a clear and interpretable physical representation of energy costs, enabling more efficient and intuitive hyperparameter tuning for improved energy efficiency. ECO introduces dedicated constraints for energy consumption and reference motion, enforced by the Lagrangian method, to achieve stable, symmetric, and energy-efficient walking for humanoid robots. We evaluated ECO against MPC, standard RL with reward shaping, and four state-of-the-art constrained RL methods. Experiments, including sim-to-sim and sim-to-real transfers on the kid-sized humanoid robot BRUCE, demonstrate that ECO significantly reduces energy consumption compared to baselines while maintaining robust walking performance. These results highlight a substantial advancement in energy-efficient humanoid locomotion. All experimental demonstrations can be found on the project website: <https://sites.google.com/view/eco-humanoid>.

Note to Practitioners—Traditional MPC and RL approaches often require extensive hyperparameter tuning and frequently result in suboptimal solutions for improving energy efficiency while maintaining stable walking performance. ECO is designed to address these challenges by reformulating energy consumption as explicit inequality constraints, providing a physically interpretable and intuitive approach to optimizing energy efficiency. This framework is particularly well-suited for applications prioritizing energy conservation and operational stability, such as surveillance, disaster response, and long-duration autonomous operations. Additionally, ECO generates emergent behaviors such

This work was supported in part by the National Natural Science Foundation of China (No. 62403064, 62403063) and Shenzhen Science and Technology Program (No. ZDCY20250901094531003). (Weidong Huang, Jingwen Zhang contributed equally to this work.) (Corresponding authors: Jingwen Zhang and Yao Su.)

Weidong Huang, Jingwen Zhang, Jiongye Li, Shibowen Zhang, Jiayang Wu, Jiayi Wang, Hangxin Liu, Yao Su are with State Key Laboratory of General Artificial Intelligence, Beijing Institute for General Artificial Intelligence (BIGAI), Beijing 100080, China (e-mails: bigeasthuang@gmail.com; zhangjingwen@bigai.ai; lijiongye@bigai.ai; zhangshibowen@bigai.ai; wujiayang@bigai.ai; wangjiayi@bigai.ai; liuhx@bigai.ai; suyao@bigai.ai).

Jiongye Li is also with Department of Automation, Tsinghua University, Beijing 100084, China.

Shibowen Zhang is also with Department of Automation, University of Science and Technology of China, Hefei 230022, China.

Jiayang Wu is also with Department of Computer Science, Harbin Institute of Technology, Harbin 150001, China.

Yaodong Yang is with Institute for Artificial Intelligence and School of Artificial Intelligence, Peking University, Beijing 100871, China (e-mail: yaodong.yang@pku.edu.cn).




Fig. 1: Comparison between the proposed constrained RL framework, ECO, with MPC and normal RL (PPO) baselines. It creates synergy between energy consumption and walking stability of the humanoid robots without requiring an extensive parameter-tuning process, outperforming both MPC and normal RL baselines.

as lighter steps and reduced body shaking, which are especially advantageous for loco-manipulation tasks by minimizing disruptions to upper-body manipulation caused by locomotion. Comparative experiments empirically offer valuable insights into constraint selection and learning setups, which may also inspire relevant ongoing research in constrained RL.

Index Terms—Humanoid and bipedal locomotion, constrained reinforcement learning, legged robots

I. INTRODUCTION

RECENT advances in humanoid robotics have enabled these robots to perform complex motions, such as walking, running, jumping, and even loco-manipulation [1–4]. Despite these achievements, energy efficiency still remains a significant bottleneck for humanoid robots. Compared to their biological counterparts [5], humanoid robots require substantially more energy to perform similar tasks, limiting their operational range, endurance, and maximum payload in real-world applications [5, 6].

To improve the energy efficiency of humanoid robots, one promising approach is to leverage reinforcement learning (RL) to optimize the control policy. For instance, Proximal Policy Optimization (PPO)-based frameworks incorporate energy-related terms into reward functions, such as minimizing joint torques, joint accelerations, and contact forces [7–9]. However,

to achieve efficient and stable locomotion, these methods often require extensive tuning on the weights of the reward terms regarding energy cost, task completion, and stability. This tuning process is non-intuitive and time-consuming, as (i) the effects of weight configurations are not straightforward, (ii) the training periods—ranging from hours to days—to evaluate a single coefficient choice even with parallel computation [10, 11]. Furthermore, conflicting objectives within the reward terms can result in suboptimal solutions or even convergence failures [12, 13]. For example, the learned policy might prioritize minimizing energy consumption at the cost of stability, resulting in unstable walking, or emphasize stability at the expense of significantly increased energy usage.

To simplify the reward tuning process for optimizing energetically efficient locomotion with RL, we propose ECO (Energy-Constrained Optimization). ECO leverages constrained RL formulation to separate conflicting objectives into distinct rewards and constraints [1, 14, 15]. This separation ensures that critical constraints, such as energy consumption limits, are strictly enforced, while task-related rewards, such as velocity tracking and stability can be optimized without extensive tuning of coefficients. Energy constraint thresholds are identified through a physically intuitive tuning process, incrementally adjusted based on direct energy criteria using a linear search algorithm. We evaluate our approach by comparing four state-of-the-art constrained RL methods and investigating various constraint settings. Our results empirically show that the proposed setup outperforms others in training stability and constraint enforcement. Through extensive experiments on both simulation and hardware platforms, we validate that ECO consistently converges to solutions that minimize energy consumption while ensuring stable and robust walking.

Our main contributions can be summarized as follows:

- 1) We propose ECO, a method designed for optimizing the energy efficiency of humanoid locomotion. Rather than balancing energy minimization with other reward terms—such as velocity tracking, motion tracking, and stability—we explicitly treat energy cost as a constraint. This approach provides a clear, interpretable physical meaning of the energy cost, enabling efficient and intuitive hyperparameter tuning focused on energy efficiency.
- 2) We investigate four constrained RL algorithms and also different constraint settings to train the policy. Through extensive simulations, we empirically find that the PPO-Lagrangian method with the current constraint setup achieves the lowest energy cost with fast and stable convergence. Additionally, instead of manually designing efficient walking, we observe that the resulting policy promotes i) extended knee movements, ii) lighter steps, and iii) reduced body shaking as shown in Fig. 1.
- 3) We carry out real-world validations on the Bruce robot. To the best of our knowledge, this is the first work to achieve energy-efficient humanoid walking using constrained RL on a real humanoid robot. Compared to MPC and PPO baselines, ECO demonstrates approximately 6 times lower energy consumption than MPC and 2.3 times lower than PPO, setting a new benchmark for energy-efficient humanoid locomotion.

We organize the remainder of the paper as follows. Sec. II reviews the literature related to this work. Sec. III introduces the preliminaries of RL and constrained RL formulations. Sec. IV presents the proposed ECO framework in detail. Secs. V and VI show the simulation and experiment result of ECO on BRUCE with comprehensive evaluations and compares with baseline methods. Finally, we discuss and conclude the paper in Secs. VII and VIII.

II. RELATED WORK

A. Energy Optimization in Legged Locomotion

Improving energy efficiency directly helps to extend the endurance of legged robots and prevent actuator overheating, thereby improving the overall performance [6]. For example, passive dynamic walkers can achieve efficient and stable walking through mechanical designs without any actuation [16] or with minimal actuation [17]. However, achieving similar efficiency in more complex legged systems, such as quadrupedal and bipedal robots, remains a challenge. As a unified Trajectory Optimization (TO) framework has been widely used for legged locomotion [18, 19], incorporating energy-related constraints tends to be a potential solution. [20] formulates the TO problem with the cost of transport (CoT) as a constraint, which quantifies the energy expenditure per unit distance. However, using TO to consider complex constraints often results in a high-dimensional and nonconvex landscape, which hinders real-time usage. To simplify the problem, energy constraints are approximated using the norm of joint torques and accelerations [2, 21, 22], or by encouraging specific bioinspired efficient behaviors, such as adaptive center of mass (CoM) height control and rolling contact of the stance leg in gait design [23]. In these methods, energy consumption is only implicitly considered, which often makes it compromised in favor of other task requirements, such as stability and robustness.

B. Learning for Legged Locomotion

Given the aforementioned issues, offline learning a locomotion policy with complex rewards can be a promising solution to balance energy-efficiency and robustness in real-time [24–26]. [27] demonstrated sim-to-real learning of multiple bipedal gaits through periodic reward composition. Similarly, [28] extended rapid motor adaptation (RMA) for bipedal locomotion to address challenges such as payloads and slippery surfaces. Other methods also integrated model-based footstep planning with model-free RL to achieve dynamic and efficient walking [29, 30]. However, to avoid disrupting the primary learning goals, energy penalties in the reward function often act as regularizers, with minimal impact on reducing excessive joint torque usage. [31] balanced energy consumption with stability through curriculum learning with symmetric mirror loss. [32] further studied different ways to incorporate symmetry, highlighting how symmetry priors can improve gait robustness and sample efficiency without overly restricting policy exploration. [7] explored efficient quadrupedal gaits with energy-centric rewards. [33–36] train control policies for humanoid locomotion using human motion

data as reference trajectories. Despite these advancements, the energy-related hyperparameters, like the reward weight or curriculum length, still present an unclear physical meaning, which intricates the weight tuning and curriculum design.

C. Constrained Reinforcement Learning

Constrained RL, or safe RL, formulates the problem as a constrained Markov decision process (CMDP) [37–40], using cost functions to decouple rewards from constraints. This allows robots to maximize rewards while keeping costs within physically meaningful thresholds. CPO [14] uses second-order optimization to enforce constraints iteratively but struggles with scalability in multi-constraint settings. PPO-Lagrangian (PPO-Lag) [41] enhances PPO with primal-dual updates for better constraint satisfaction and performance. Constraint manifold is defined for safe exploration leading to actions that satisfy corresponding constraints in [42, 43]. However, efficient constraint manifold computation for legged robots with a floating base remains challenging. Constraint-rectified Policy Optimization (CRPO) [44] offers a primal-based framework well-suited for multi-constraint problems, while Interior-point Policy Optimization (IPO) [45] employs log-barrier penalties for feasibility enforcement. [15] extended IPO with a teacher-student framework, transferring most reward terms into constraints to avoid reward scale tuning, and demonstrated robust performance on legged robots. Penalized Proximal Policy Optimization (P3O) [46] uses clamping penalty functions for constraint enforcement and has shown success in reducing violations (*e.g.*, self-collision, joint speed limits) in quadruped locomotion [47]. Despite these advancements, these methods have only been validated on quadruped robots. Applying them to humanoid robots, especially when combining energy consumption with other constraints (*e.g.*, self-collision), can overly restrict the search space. Given the smaller convergence region of humanoid robots compared to quadrupeds, rethinking how to select and design constraints and how to select an appropriate optimization method within the constrained RL framework to generate efficient and stable walking performance of humanoid robots remains a critical challenge.

III. PRELIMINARIES & PROBLEM STATEMENT

A. Markov Decision Processes

RL is typically modeled as a Markov Decision Process (MDP) [37]:

$$\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathbb{P}, R, \mu, \gamma), \quad (1)$$

where \mathcal{S} and \mathcal{A} denote the state and action spaces, respectively. The functions $\mathbb{P}(\mathbf{s}'|\mathbf{s}, \mathbf{a})$ and $R(\mathbf{s}'|\mathbf{s}, \mathbf{a})$ represents the transition probability and corresponding reward of moving from state \mathbf{s} to \mathbf{s}' under action \mathbf{a} . $\mu(\cdot)$ represents the initial state distribution, and γ is the discount factor.

A stationary policy π_{θ} governs the action probability $\pi_{\theta}(\mathbf{a}|\mathbf{s})$ for a given state \mathbf{s} , with all policies parameterized by θ . The corresponding infinite-horizon reward is defined as

$$J^R(\pi_{\theta}) = \mathbb{E}_{\mathbf{s} \sim \mu}^{\pi} \left[\sum_{t=0}^{\infty} \gamma^t R(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t) \right]. \quad (2)$$

The objective of the MDP is to maximize:

$$\max_{\pi_{\theta} \in \Pi_{\theta}} J^R(\pi_{\theta}), \quad (3)$$

where Π_{θ} is the neural network parameter space.

B. Constrained Markov Decision Processes

A CMDP extends the standard MDP framework by incorporating a set of cost functions, referred to as **general constraints**. This set is defined as $\mathcal{C} = (C_i, b_i)_{i=1}^m$, where each cost function $C_i : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is associated with a corresponding threshold b_i .

The objective of the CMDP is to identify the policy that maximizes the reward while satisfying all specified constraints:

$$\max_{\pi_{\theta} \in \Pi_{\mathcal{C}}} J^R(\pi_{\theta}), \quad (4)$$

where $\Pi_{\mathcal{C}} = \Pi_{\theta} \cap \{\pi_{\theta} \mid J^{C_i}(\pi_{\theta}) \leq b_i, \forall i = 1, \dots, m\}$ represents the set of feasible policies that satisfy all constraints and $J^{C_i}(\pi_{\theta})$ is the expected value of a cost function under a policy π . These constraints can take various forms, such as a discounted cumulative sum, average sum, *etc.* [37].

Discounted sum constraints are formulated to ensure that a specified quantity remains below a given threshold over the long term, accounting for cumulative effects rather than instantaneous values at each time step. They are defined as:

$$J^{C_i}(\pi_{\theta}) = \mathbb{E} \left[\sum_{t=0}^{T-1} \gamma^t C_i(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t) \right] \leq b_i, \quad (5)$$

where T is the episode length. This type of constraint is particularly well-suited for maintaining long-term requirements in robotic operations, such as limiting energy consumption over a certain period.

Average sum constraints ensure that the mean value of a specific variable remains below a given threshold at each time step. Formally, these constraints are expressed as:

$$J^{C_i}(\pi_{\theta}) = \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} C_i(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t) \right] \leq b_i. \quad (6)$$

For instance, these constraints can be utilized to enforce specific motion styles, such as symmetry during walking.

C. Constraints RL Methods

1) **PPO-Lag**: PPO-Lag [41] is one of the most widely used methods in constrained RL. The Lagrangian function is:

$$L(\theta, \lambda_1, \dots, \lambda_m) = J^R(\pi_{\theta}) - \sum_{i=1}^m \lambda_i (J^{C_i}(\pi_{\theta}) - b_i), \quad (7)$$

where λ_i is the Lagrange multipliers. Using this approach, the original constrained optimization problem in Eq. (4) is reformulated as an unconstrained optimization problem:

$$\max_{\pi_{\theta} \in \Pi_{\theta}} \min_{\lambda_1, \dots, \lambda_m \geq 0} L(\theta, \lambda_1, \dots, \lambda_m). \quad (8)$$

This optimization process alternates between updating the primal variable π_{θ} and the dual variables λ_i , ultimately yielding an approximate solution to the Eq. (4).




Fig. 2: **Overview of the training and deployment process in proposed ECO framework.** The policy network, taking velocity commands and proprioception data as inputs, outputs desired joint positions at 100Hz to a PD controller, which updates torque commands at 1kHz. The reward critic is trained with privileged observations. The simulator provides the reward, energy cost, and symmetry cost, which are used to compute the reward and cost returns. The policy is then updated using the Lagrangian formulation in Eq. (20) to balance rewards and costs. The trained policy is directly deployed to the real world.

2) **CRPO**: CRPO [44] alternates between maximizing the reward and minimizing the cost, prioritizing the reduction of costs when constraints are violated. In cases where multiple constraints are violated, a single constraint is randomly selected for optimization:

$$L_{\text{CRPO}} = \mathbb{1}_{n_k=0} \cdot J^R(\pi_{\theta}) - \sum_{i=1}^{n_k} \nu_i \cdot J^{C_i}(\pi_{\theta}), \quad (9)$$

where $n_k = \sum_i^m \mathbb{1}_{J^{C_i} > b_i}$ is the number of violated constraints, $\nu \in \mathbb{R}^{n_k}$ is a random one-hot vector used to select a violated constraint to update at each iteration.

3) **IPO**: IPO [45] transforms the constrained RL problem into an unconstrained one by employing a logarithmic barrier function, defined as follows:

$$L_{\text{IPO}} = J^R(\pi_{\theta}) - \sum_{i=1}^m \kappa_i^{\text{IPO}} \cdot \log(J^{C_i}(\pi_{\theta}) - b_i), \quad (10)$$

where $\kappa_i^{\text{IPO}} > 0$ is the fixed hyperparameter that controls the steepness of the logarithmic function.

4) **P3O**: P3O [46] extends the PPO framework by incorporating penalties for constraint violations into the objective function. The objective is defined as:

$$L_{\text{P3O}} = J^R(\pi_{\theta}) - \sum_{i=1}^m \kappa_i^{\text{P3O}} \cdot \max\{J^{C_i}(\pi_{\theta}) + \epsilon_i, 0\}, \quad (11)$$

where κ_i^{P3O} represents the fixed weight for each constrain.

$$\epsilon_i = \mathbb{E}_{(\mathbf{s}, \mathbf{a}) \sim \text{buffer}} \left[\sum_{t=0}^T \gamma^t C_i(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t) \right] - b_i, \quad (12)$$

represents the amount by which the actual cost exceeds the constraint threshold b_i , as calculated using data collected from prior interactions with the environment.

In this paper, we compare these methods and ultimately select PPO-Lag as the constrained optimization approach for

ECO. In Sec. IV, we will provide detailed implementation insights and explain the rationale behind this choice. In Sec. V, we will compare their performance to validate this selection.

IV. PROPOSED ECO FRAMEWORK

In this section, we describe the proposed ECO framework, which aims to minimize energy consumption in humanoid locomotion while ensuring stable and natural leg movements. Both the control policy and reward critic are implemented as fully connected neural networks. As illustrated in Fig. 2, the control policy takes the robot's proprioception and velocity commands as input and outputs the desired joint positions. We train the policy using PPO-Lagrangian [48] in IsaacGym and deploy it directly on the physical robot. The following subsections introduce the state space, action space, constraints settings, reward settings, and the formulation for constraint policy optimization.

A. State Space

The input to the policy network consists of an observation history comprising K_f consecutive frames, which include the robot's proprioceptive states and velocity commands, with a dimensionality of \mathcal{S} . Specifically, \mathcal{S} includes velocity commands $(v_x^{\text{cmd}}, v_y^{\text{cmd}}, \omega_{\text{yaw}}^{\text{cmd}}) \in \mathbb{R}^3$, clock inputs $(\sin(t), \cos(t)) \in \mathbb{R}^2$, joint positions $\mathbf{q} \in \mathbb{R}^n$, joint velocities $\dot{\mathbf{q}} \in \mathbb{R}^n$, body angular velocities $\boldsymbol{\omega}_{\text{ang}} \in \mathbb{R}^3$, last action $\mathbf{a}_{t-1} \in \mathbb{R}^n$, and the XY components of body Euler angles $\boldsymbol{\omega}_{\text{ang}} \in \mathbb{R}^2$, where n is number of leg Degree of Freedoms (DoFs). The body yaw angle is excluded due to its drift on the real hardware.

The state space of our reward critic includes both the inputs used by the policy network and additional privileged information, such as body linear velocity, body yaw angle, external push forces and torques, friction coefficients, body mass, phase identifiers, and contact indicators. These indicators specify the

swing and stance phases of each foot, as determined by the contact scheduler.

B. Action Space

In this work, we actively control leg joints while the arm joints are fixed to nominal angles. The action from the policy, represented as $\mathbf{a} \in \mathbb{R}^n$, is defined as the deviations from the nominal joint positions $\mathbf{q}^{\text{nominal}}$ [29, 49]. This action is fed into the PD controller to convert the action into the desired torque command using the following equation:

$$\hat{\boldsymbol{\tau}} = \mathbf{K}_p(\mathbf{a} + \mathbf{q}^{\text{nominal}} - \mathbf{q}) + \mathbf{K}_d(\mathbf{0} - \dot{\mathbf{q}}), \quad (13)$$

where $\mathbf{q}, \dot{\mathbf{q}}$ denote the current joint position and joint velocity.

C. Constraints

1) **Constraints Selection Consideration:** Since minimizing energy consumption is the main focus of this paper, a naive choice is to mix it with other objectives, such as adherence to task commands, behavior stability, and disturbance rejection. However, regarding energy consumption, we observed that adjusting the reward weights alone for PPO does not lead to convergence within the feasible domain (shown in Sec. V-D1). Consequently, we select energy consumption as a constraint in this optimization process. Additionally, we integrate reference motion (using Mirror Loss as the reference in this paper) as a constraint to encourage natural motions.

Among prior works [15, 47], additional reward terms are often placed within constraints for quadrupedal robots. In Sec. V-D6, we investigate various constraint settings including the constraints used in [15, 47], and find that the current approach struggles with stable convergence as the number of constraints increases. This may be due to humanoid robots having a much smaller constrained feasible domain than quadrupeds, which makes convergence more challenging in the presence of multiple constraints. Therefore, in this work, we focus on two primary constraints: energy minimization and reference motion tracking for humanoid robots.

2) **Energy Constraint:** Due to the challenges in accurately measuring link inertia and quantifying friction and heat dissipation losses, the overall energy consumption of a robot cannot be directly measured. Therefore, in this work, the motor energy consumption, which constitutes the primary part of the total energy consumption, is selected as the evaluation metric. It can be calculated by integrating the absolute power of the n motors. Since the timestep dt is fixed, the energy cost function is defined as

$$C_1(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t) = \sum_{j=1}^n |\boldsymbol{\tau}_t^j \dot{\mathbf{q}}_t^j| \quad (14)$$

to exclude dt . Here, $\boldsymbol{\tau}_t^j$ and $\dot{\mathbf{q}}_t^j$ denote the torque and velocity of the j -th joint at time t . Using the discounted sum constraint from Eq. (5), the energy constraint is formulated as:

$$J^{C_1}(\boldsymbol{\pi}_\theta) = \mathbb{E} \left[\sum_{t=0}^T \gamma^t \sum_{j=1}^n |\boldsymbol{\tau}_t^j \dot{\mathbf{q}}_t^j| \right] \leq b_1, \quad (15)$$

where b_1 is the energy threshold, estimated based on the robot's speed range and walking duration.

3) **Reference Motion Constraint:** To enforce stability and mirror symmetry, we define a mirror loss term following [31] as the cost function:

$$C_2(\mathbf{s}'|\mathbf{s}, \mathbf{a}) = \|sg(\boldsymbol{\pi}_\theta(\mathbf{s}_i)) - \Psi_a(\boldsymbol{\pi}_\theta(\Psi_o(\mathbf{s}_i)))\|^2, \quad (16)$$

where $sg(\cdot)$ is the stop gradient operator, $\Psi_a(\cdot)$ and $\Psi_o(\cdot)$ are the mirroring functions for actions and states, respectively, and $\boldsymbol{\pi}_\theta(\mathbf{s}_i)$ is the mean action of the stochastic policy conditioned on the state \mathbf{s}_i . This constraint ensures that when the observation input to the policy is mirrored, the corresponding action output is also mirrored. Practically, if the policy learns to lift the left leg at the appropriate moment, this constraint also enforces it to lift the right leg at the corresponding moment, and vice versa. This reference constraint promotes more stable and symmetric behavior within the policy. Such symmetry fosters balanced behavior without the need to prescribe motion trajectories explicitly, thereby reducing the necessity to define trajectory parameters manually.

We apply the average constraint from Eq. (6) to this cost function:

$$J^{C_2}(\boldsymbol{\pi}_\theta) = \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^T \|\boldsymbol{\pi}_\theta(\mathbf{s}_t) - \Psi_a(\boldsymbol{\pi}_\theta(\Psi_o(\mathbf{s}_t)))\|^2 \right] \leq b_2, \quad (17)$$

where b_2 is the constraint threshold of reference motion.

D. Rewards

The aspect of reward design is not the main focus of this paper, therefore we adopt a reward configuration similar to the Humanoid-Gym [9]. Since reference motion and energy constraints are directly integrated into the cost function, we have omitted the reference joint position tracking reward, joint velocity penalty reward, and joint torque penalty reward from the original setup. The total reward function is defined as

$$R(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t) = \sum_i \mu_i \cdot r_i, \quad (18)$$

where the rewards r_i and their corresponding scales μ_i will be detailed in Sec. V.

E. Constrained Policy Optimization

To solve the constrained optimization for energy-efficient humanoid walking, we employ the Lagrangian method to transform the original constrained objective into an unconstrained form, as introduced in Sec. III. Having two designed constraints, we formulate the problem as follows:

$$\max_{\boldsymbol{\pi}_\theta \in \Pi_\theta} \min_{\lambda_1, \lambda_2 \geq 0} J^R(\boldsymbol{\pi}_\theta) - \sum_{i=1}^2 \lambda_i (J^{C_i}(\boldsymbol{\pi}_\theta) - b_i). \quad (19)$$

Here, $J^{C_1}(\boldsymbol{\pi}_\theta)$ and $J^{C_2}(\boldsymbol{\pi}_\theta)$ are cost returns defined by Eqs. (15) and (17). Solving Eq. (19) involves iteratively updating the policy parameters and the Lagrange multipliers. The update steps alternate between improving the policy based on both reward and cost gradients, and adjusting λ_i to enforce constraints. Formally:

$$\begin{aligned} \boldsymbol{\pi}_{\theta^{k+1}} &= \boldsymbol{\pi}_{\theta^k} + \alpha(k) \left[\nabla_{\boldsymbol{\theta}^k} J^R(\boldsymbol{\pi}_{\theta^k}) - \sum_{i=1}^2 \lambda_i^k \nabla_{\boldsymbol{\theta}^k} J^{C_i}(\boldsymbol{\pi}_{\theta^k}) \right], \\ \lambda_i^{k+1} &= \max\{0, \lambda_i^k - \beta_i(k) (J^{C_i}(\boldsymbol{\pi}_{\theta^k}) - b_i)\}. \end{aligned} \quad (20)$$

Here, $\alpha(k)$ and $\beta_i(k)$ are learning rates computed by Adam optimizer [50] at the training iteration k . The projection $\max\{0, x\}$ ensures that each Lagrange multiplier remains non-negative. We integrate this Lagrangian approach into the PPO algorithm [51] within an Actor-Critic framework. The policy is optimized to maximize a combined objective:

$$L_{\text{PPO-Lag}} = L_R - r_t(\boldsymbol{\theta}) \sum_{i=1}^2 \lambda_i J^{C_i}(\pi_{\boldsymbol{\theta}}), \quad (21)$$

where L_R is the clipped surrogate objective defined by PPO:

$$L_R = \min\{r_t(\boldsymbol{\theta}) A^{\pi_{\text{old}}}(\mathbf{s}_{\leq t}, \mathbf{a}_t), \text{clip}(r_t(\boldsymbol{\theta}), c_1, c_2) A^{\pi_{\text{old}}}(\mathbf{s}_{\leq t}, \mathbf{a}_t)\}, \quad (22)$$

with the probability ratio

$$r_t(\boldsymbol{\theta}) = \frac{\pi_{\boldsymbol{\theta}}(\mathbf{a}_t | \mathbf{s}_{\leq t})}{\pi_{\text{old}}(\mathbf{a}_t | \mathbf{s}_{\leq t})}. \quad (23)$$

In this formulation, π_{old} denotes the behavior policy used for data collection, $A^{\pi_{\text{old}}}$ is the advantage computed via GAE [52], and c_1, c_2 define the PPO clipping range.

Similarly, within the PPO framework, the final objectives of IPO, P3O, and CRPO are as follows:

$$L_{\text{IPO}} = L_R - r_t(\boldsymbol{\theta}) \cdot \sum_{i=1}^2 \kappa_i^{\text{IPO}} \cdot \log(J^{C_i}(\pi_{\boldsymbol{\theta}}) - b_i),$$

$$L_{\text{P3O}} = L_R - r_t(\boldsymbol{\theta}) \cdot \sum_{i=1}^2 \kappa_i^{\text{P3O}} \cdot \max\{J^{C_i}(\pi_{\boldsymbol{\theta}}) + \epsilon_i, 0\}, \quad (24)$$

$$L_{\text{CRPO}} = \mathbb{1}_{n_k=0} \cdot L_R - r_t(\boldsymbol{\theta}) \cdot \sum_i^2 \nu_i \cdot J^{C_i}(\pi_{\boldsymbol{\theta}}).$$

Among the algorithms compared, P3O (with fixed $\kappa_i^{\text{P3O}} = 1.0$) and PPO-Lag (with initial $\lambda_i = 0$, $\beta_i = 1e-3$) required minimal parameter tuning and achieved low constraint violations. While P3O demonstrated comparable energy performance, PPO-Lag exhibited better convergence speed and stability, thereby selected for implementing the constrained policy optimization approach in ECO.

The reward critic is trained with the following loss [53]:

$$L_{\text{Reward Critic}} = \max\{\|V(\mathbf{s}_t) - R_t\|_2, \|V_{\text{target}}(\mathbf{s}_t) + \text{clip}(V_{\text{target}}(\mathbf{s}_t) - V(\mathbf{s}_t), c_1, c_2) - R_t\|_2\}, \quad (25)$$

where R_t is the cumulative return at time t , $V(\mathbf{s}_t)$ denotes the value estimated by the critic at state \mathbf{s}_t , and $V_{\text{target}}(\mathbf{s}_t)$ is the value approximated by the target critic at state \mathbf{s}_t . Unlike previous studies [41, 48] that utilize a cost critic, we estimate cost return for energy using the Monte Carlo approach [54], which achieves comparable performance while requiring fewer network parameters and minimal design choices.

V. SIMULATION

A. Robot Platform

We use the kid-sized humanoid robot BRUCE [3, 55, 56] as an example to demonstrate the performance of the proposed ECO framework. BRUCE is 70cm tall, weighs 4.8kg, and has a total of 16 DoFs, with 3 per arm and 5 per leg. All leg DoFs are driven by liquid-cooled proprioceptive actuators with peak torque 10.5Nm, while the arm DoFs are driven by Dynamixel servo motors. The PD controller gains for both legs are set as $\mathbf{K}_p = \text{diag}(7, 10, 7, 10, 1.5)$ and $\mathbf{K}_d = \text{diag}(0.2, 0.4, 0.2, 0.4, 0.08)$ for the hip yaw, hip pitch, hip roll, knee pitch, and ankle pitch joints, respectively.

B. Learning Setups

1) **Training Setups:** The policy training is conducted in the IsaacGym [57] with parallelization across 8192 environments. The utilized rewards and scales are summarized in Tab. I. The constraint policy optimization of ECO utilizes the PPO-Lagrangian [41] and the hyperparameters are shown in Tab. III. We employ domain randomization [58] to enhance transfer to unseen scenarios, with specific parameters listed in Tab. II. During training, two types of external perturbations are applied to enhance the policy robustness: force disturbances every 2s, lasting 0.001s with magnitudes randomly sampled between $[-100, 100]N$, and velocity impulses every 4s, also lasting

TABLE I: **Summary of additional reward functions.** The tracking error metric is defined as: $\phi(e, w) := \exp(-w \cdot \|e\|^2)$, where e represents the tracking error, and w is the associated weight. Δv_{lin} and $\Delta \omega_{\text{ang}}$ represent the linear and angular velocity errors; $\mathbb{1}_{(\cdot)}$ is the indicator function for the logical condition (\cdot), for example, $\mathbb{1}_{\text{stand phase}} = 1$ when the leg is in stand phase. t_{air} is the duration that the foot remains in the air during locomotion; \dot{P}^f is the velocity of the feet; $q_{\text{yaw/roll}}$ is the yaw and roll joint position; q_0 is the default joint position; P_z^b is the base height; Δh_f is the deviation in foot height; h_t is the target foot height during the swing phase; \sum represents the summation across all legs. The base height target is set to 0.45m, the maximum contact force F_{max} is set to 50N, and the target foot height h_t is 0.03m.

Reward	Equation (r_i)	Scale (μ_i)
Tracking Linear Velocity	$\phi(\Delta v_{\text{lin}}, 10)$	1.2
Tracking Angular Velocity	$\phi(\Delta \omega_{\text{ang}}, 10)$	1.1
Foot Air Time	$\sum \mathbb{1}_{\text{contact}} \cdot \min(t_{\text{air}}, 0.5)$	1.0
Foot Contact Velocity	$\sum \sqrt{\ \dot{P}^f\ } \cdot \mathbb{1}_{\text{contact}}$	-0.05
Foot Clearance	$\sum \mathbb{1}_{ \Delta h_f - h_t < 0.01} \cdot \mathbb{1}_{\text{swing phase}}$	1.0
Foot Contact Number	$\frac{1}{N} \sum (\mathbb{1}_{\text{contact}} = \mathbb{1}_{\text{stand phase}})$	1.2
Base Orientation	$\frac{1}{2} (\phi(\text{euler angles}, 10) + \phi(\text{gravity vector}, 20))$	1.0
Foot Contact Forces	$\sum \max(0, \ F_c\ - F_{\text{max}})$	-0.01
Default Joint Pos	$\phi(q_{\text{yaw/roll}} - q_0, 100)$	0.5
Tracking Base Height	$\phi(P_z^b - 0.45, 100)$	0.2
Self-collision Penalty	$\sum (\mathbb{1}_{\text{self-collision}})$	-1.0
Action Smoothness	$\ \mathbf{a}_t - \mathbf{a}_{t-1}\ ^2 + \ \mathbf{a}_t - 2\mathbf{a}_{t-1} + \mathbf{a}_{t-2}\ ^2 + 0.05\ \mathbf{a}_t\ $	-0.002

TABLE II: **Overview of domain randomization.** The domain randomization terms and the associated parameter ranges are listed. Additive randomization increments the parameter by a value within the specified range while scaling randomization adjusts it by a multiplicative factor from the same range.

Parameter	Unit	Range	Operator	Type
Payload	kg	[-0.5, 0.5]	additive	Uniform
Body COM Displacement	m	[-0.05, 0.05]	additive	Uniform
Floor Friction	-	[0.1, 2.0]	-	Uniform
Restitution	-	[0.0, 0.5]	additive	Uniform
Motor Strength	%	[0.9, 1.1]	scaling	Uniform
Joint Friction	%	[0.02, 0.05]	scaling	Uniform
Joint armature	%	[0.0, 0.01]	additive	Uniform
Action Delay	ms	[0, 10]	-	Uniform
Kp Factors	%	[0.9, 1.1]	scaling	Uniform
Kd Factor	%	[0.9, 1.1]	scaling	Uniform
Initial Euler Angle Noise	rad	[-0.1, 0.1]	additive	Uniform
Joint Position Noise	rad	[-0.03, 0.03]	additive	Gaussian (1σ)
Joint Velocity Noise	rad/s	[-0.6, 0.6]	additive	Gaussian (1σ)
Angular Velocity Noise	rad/s	[-0.06, 0.06]	additive	Gaussian (1σ)
Euler Angle Noise	rad	[-0.018, 0.018]	additive	Gaussian (1σ)

TABLE III: **Hyperparameters in the training process.**

Parameter	Value
Number of Environments	8192
Number Training Epochs	2
Batch size	8192×96
Episode Length	24s
Discount Factor of Reward	0.994
Discount Factor of Energy Cost	0.90
PPO Clip Coefficient c_1, c_2	0.8, 1.2
Frame Stack of Single Observation K_f	15
Frame Stack of Single Privileged Observation	3
Number of Single Observation S	40
Number of Single Privileged Observation	55
Initial Lagrangian Multipliers λ_1, λ_2	0.0, 0.0
Learning Rate of Lagrangian Multiplier β_1, β_2	0.001, 0.001
Cost Threshold b_1 (0.1m/s)	$60J \approx 2.5W \times 24s$
Cost Threshold b_1 (0.15m/s)	$70J \approx 2.9W \times 24s$
Cost Threshold b_1 (0.2m/s)	$80J \approx 3.3W \times 24s$
Cost Threshold b_2	0.05

0.001s, with linear velocities up to 0.2m/s and angular velocities up to 0.4rad/s. The simulation is executed at a frequency of 1kHz, with a control frequency of 100Hz. The training process is carried out on a system equipped with an Intel Core i9-13900KFN processor and NVIDIA GeForce RTX 4090, and each training run consists of 2000 policy iterations. Convergence is typically achieved within approximately 4h of wall-clock time. We use the policy trained after 10h for sim-to-real experiments.

2) **Constraint Thresholds Search:** Since the PPO baseline policy described in Sec. V-C converges to a suboptimal solution, it serves as a useful initial reference for selecting constraint thresholds. For the reference motion constraint, we empirically observed that the mirror-symmetry loss value obtained from the PPO baseline ensures robust and symmetric walking behavior. Accordingly, we set the threshold b_2 in Eq. (17) as 0.05 across all experiments. For the energy constraint threshold, we adopt a linear search strategy [15] to determine appropriate values, further details are provided in Sec. V-D4. Each threshold search takes around 4 hours of wall-clock time depending on the chosen step size. Additionally, as summarized in Tab. III, energy consumption thresholds may vary with different target velocities.

C. Baselines

For comparison, five baseline approaches are chosen to demonstrate that our proposed ECO framework significantly improves energy efficiency without degenerating the performance of humanoid robots in walking tasks. The **MPC baseline** uses the state-of-the-art model-based hierarchical method which combines a high-level MPC planner with a simplified model and a low-level whole-body controller with the full-body dynamics [19]. Humanoid-Gym [9] is used as the **model-free RL baseline (PPO)** for training end-to-end locomotion policies in humanoid robots. To facilitate a rigorous comparison with the ECO model, the PPO baseline has been augmented with an energy penalty reward [7], which serves as the constraint within the ECO framework. We selected scales of energy penalty reward $\mu_e \in \{-0.001, -0.01, -0.1, -0.03, -0.05, -0.08, -0.015, -0.02, -0.025\}$, while all other rewards and scales are aligned with those of ECO as detailed in Tab. I. Additionally, a mirror loss is included in the PPO loss function to encourage symmetry, with a fixed weight of 0.1. Three constrained RL baseline methods are chosen as introduced in Eq. (24). For IPO and P3O, we fine-tune the hyperparameters as follows: $\kappa_1^{P3O} = 1.0$, $\kappa_2^{P3O} = 1.0$, $\kappa_1^{IPO} \in \{1.0, 10.0\}$, and $\kappa_2^{IPO} = 1.0$.

D. Simulation Results

1) **Comparing Constraint RL Algorithms:** We compare the proposed ECO framework against PPO, IPO, P3O, and CRPO under a 0.1m/s linear velocity tracking task, subject to energy and reference motion constraints. To assess the performance, we evaluate four key metrics, each capturing a distinct aspect of the task: (1) Normalized total reward \hat{J}^R — reflects overall task performance; (2) Normalized episode length T_{ep} — measures stability and robustness; (3) Mean energy cost \hat{J}^{C_1} — quantifies energy efficiency; (4) Mean mirror cost \hat{J}^{C_2} — evaluates motion symmetry. The metrics are formally defined as:

$$\hat{J}^R = \frac{\frac{1}{E} \sum_{i=1}^E \sum_{t=0}^T r^{t,i} - R_{\min}}{R_{\max} - R_{\min}}, \quad (26)$$

$$T_{ep} = \frac{\frac{1}{E} \sum_{i=1}^E T_{ep}^i - T_{ep}^{\max}}{T_{ep}^{\max} - T_{ep}^{\min}}, \quad (27)$$

$$\hat{J}^{C_1} = \frac{1}{E} \sum_{i=1}^E \sum_{t=0}^T C_1(\mathbf{s}_t^i, \mathbf{a}_t^i), \quad (28)$$

$$\hat{J}^{C_2} = \frac{1}{B} \sum_{i=1}^B C_2(\mathbf{s}_i, \mathbf{a}_i), \quad (29)$$

where $r^{t,i}$ is the reward at time t of episode i , and R_{\min}, R_{\max} are the minimum and maximum cumulative rewards observed across all algorithms, T_{ep}^i is the length of episode i , and $T_{ep}^{\min}, T_{ep}^{\max}$ are the minimum and maximum episode lengths among all methods, $(\mathbf{s}_i, \mathbf{a}_i)$ is the i -th state-action sample in a batch and B is the batch size. We enforce the average-sum constraint over each gradient batch as in [59], using the same batch size across all algorithms.

As shown in Fig. 3, ECO (initialized with $\lambda_i = 0, \beta_i = 10^{-3}$) consistently converges to the target energy threshold of




Fig. 3: **Comparison of training metrics for ECO, P3O, IPO, and CRPO.** The energy consumption and mirror reference motion thresholds are set at $60J$ and 0.05 , respectively, as indicated by the **black dashed lines** in (a) and (b). Results averaged over 10 random seeds.




Fig. 4: **Comparison of training metrics for ECO and the PPO.** The energy consumption and mirror reference motion thresholds are set at $60J$ and 0.05 , respectively, as indicated by the **black dashed lines** in (a) and (b). Results averaged over 10 random seeds.

$60J$ while maintaining a normalized episode length close to 1.0, demonstrating both energy efficiency and training stability. In comparison, P3O satisfies both the energy and reference motion constraints but requires more iterations to converge. IPO exhibits sensitivity to the penalty coefficient κ_1^{IPO} : with $\kappa_1^{\text{IPO}} = 1.0$, it fails to satisfy the energy constraint; increasing it to $\kappa_1^{\text{IPO}} = 10.0$ enforces the constraint but leads to reduced episode length due to overly conservative behavior. CRPO,

while capable of satisfying the constraints, exhibits unstable convergence, likely due to frequent policy shifts between optimizing for reward and constraint cost.

As shown in Fig. 4, we explore three orders of magnitude for the energy reward coefficient in PPO (μ_e), but none yields a satisfactory balance between energy efficiency and stability. Lower coefficients failed to reduce energy consumption, while higher ones caused instability and frequent falls. These results




Fig. 5: **Sim-to-Sim transfer results.** For visual clarity, single-run deployment curves are shown. (a) Visual comparison of walking stability across simulators; (b) Ankle height consistency across simulators; (c) Motor energy consumption at $0.1m/s$ over $10s$ in Gazebo; (d) Body velocity tracking in MuJoCo in different speed commands.

highlight the difficulty of manually tuning reward weights in multi-objective PPO. In contrast, ECO directly enforces energy and motion constraints, avoiding such tuning and achieving robust, energy-efficient behavior.

Regarding the reference motion constraint (Fig. 3(b)), IPO, P3O, and CRPO rapidly converge to the predefined threshold of 0.05, though their resulting behaviors are more conservative than those of ECO. In contrast, PPO with various energy coefficients converges to different levels of mirror cost (Fig. 4(b)), often yielding asymmetric and unstable gaits. These results highlight that constrained RL offers a more principled and efficient framework for aligning policy behaviors with target objectives without heavy hyperparameter tuning.

Finally, we note that PPO, P3O, and IPO rely on fixed constraint coefficients (μ_e , κ^{P3O} , κ^{IPO}), and CRPO lacks constraint-specific parameters altogether. These approaches are therefore prone to convergence issues or require careful manual tuning. In contrast, ECO leverages PPO-Lag to dynamically update the Lagrange multipliers λ_i using the Adam optimizer, allowing adaptive scaling the impact of constraints during training. This dynamic adjustment significantly improves convergence stability and highlights PPO-Lag as a more effective strategy for optimizing energy-efficient humanoid locomotion.

2) **Sim-to-Sim Transfer:** To evaluate the robustness and generalizability of the learned policy, we transferred it to two additional simulation environments, MuJoCo [60] and Gazebo [61], to assess ECO’s sim-to-sim performance and benchmark it against baseline methods. As shown in Fig. 5(a), our method achieves stable walking across three simulation environments. Fig. 5(b) and (c) demonstrate that ECO consistently maintains a similar gait pattern during cross-simulator transfer and reduces energy usage in Gazebo while BRUCE walks at $0.1m/s$ over a $10s$ period. Fig. 5(d) further shows that our method accurately tracks linear velocity commands at three different target speeds ($0.1m/s$, $0.15m/s$, $0.2m/s$) in MuJoCo. We report summary statistics over 10 runs in Tab. IV.

Results show that MPC incurs the highest energy cost and fails to track higher target speeds (e.g., $0.2m/s$), likely due to its QP solver being unable to resolve contact dynamics in finite time. PPO exhibits better performance than MPC but with moderate energy consumption. ECO, in contrast, not only achieves accurate velocity tracking but also maintains significantly lower energy consumption—approximately 3 times lower than MPC and 1.4 times lower than PPO.

3) **Disturbance Rejection Test:** To verify that the walking robustness has not been sacrificed in the learned policy, we conduct disturbance rejection tests on BRUCE while walking




Fig. 6: **Benchmark of disturbance rejection.** BRUCE was subjected to pushes of varying durations from 8 different directions (0° , 45° , 90° , 135° , 180° , 225° , 270° , 315°). Results are averaged over 10 runs.

TABLE IV: **Velocity tracking and energy consumption across simulators.** Aggregated over 10 runs, each lasting 10 seconds.

Method	Target: (0.1m/s, 2.5W)		Target: (0.15m/s, 2.9W)		Target: (0.2m/s, 3.3W)	
	Speed (m/s)	Energy (W)	Speed (m/s)	Energy (W)	Speed (m/s)	Energy (W)
PPO (Isaacgym)	0.1129 ± 0.00071	3.1301 ± 0.0165	0.1534 ± 0.0030	4.0965 ± 0.0591	0.2141 ± 0.0043	4.6311 ± 0.0666
Ours (Isaacgym)	0.1026 ± 0.00279	2.1952 ± 0.0458	0.1617 ± 0.00188	2.2401 ± 0.0239	0.2131 ± 0.0041	2.8474 ± 0.0641
PPO (Mujoco)	0.1092 ± 0.0005	2.8150 ± 0.00895	0.1567 ± 0.00042	3.7186 ± 0.00957	0.2028 ± 0.00089	4.0692 ± 0.0162
Ours (Mujoco)	0.10033 ± 0.00064	2.00009 ± 0.0066	0.1572 ± 0.00022	2.7098 ± 0.00561	0.2038 ± 0.00064	2.8497 ± 0.00838
PPO (Gazebo)	0.1062 ± 0.0011	3.1508 ± 0.0154	0.1528 ± 0.00092	4.3922 ± 0.0181	0.2027 ± 0.0037	5.1826 ± 0.0434
MPC (Gazebo)*	0.1075 ± 0.0028	9.2693 ± 0.1604	0.1322 ± 0.0047	9.9155 ± 0.3843	0.1523 ± 0.0057	10.8626 ± 0.3497
Ours (Gazebo)	0.0930 ± 0.0029	2.3653 ± 0.0201	0.1542 ± 0.0013	2.9737 ± 0.0133	0.2159 ± 0.0035	3.5067 ± 0.0559

* MPC baseline method is exclusively implemented in Gazebo so that its performances in Isaacgym and Mujoco are disregarded.

in Gazebo. External forces are applied at the origin of the coordinate frame attached to BRUCE’s body link. The timing is chosen as vulnerable transitions from a double-stance phase to lifting the left foot or the right foot. Since the response to perturbations varies with the direction in these two cases, we record the maximum force BRUCE could resist for both timings and use the smaller of the two values for analysis. As shown in Fig. 6, the proposed ECO framework demonstrates greater robustness across all directions, especially in the sagittal direction, compared to the baseline MPC, and shows comparable robustness to PPO. Although ECO has not encountered pushes with a duration longer than $0.001s$ in the training stage, ECO still performs well under longer pushes, indicating that our method has not overfitted to improve energy efficiency in compromise of robustness. Besides controlled pushes, two qualitative disturbance experiments in Mujoco are conducted to evaluate robustness under more diverse and unpredictable scenarios: (1) different links of BRUCE are manually dragged at random timings with random magnitudes, producing varied disturbance directions and intensities; (2) Balls are thrown to BRUCE with random velocities to create different impulsive contacts. As shown in the accompanying video, ECO successfully maintains stable walking and recovers from these random perturbations, further demonstrating that the learned policy generalizes beyond the specific disturbance patterns used in training.

4) **Energy Constraint Threshold Analysis:** We observe that the baseline method PPO converges (achieving a normalized episode length of 0.9) within 1500 training iterations while consuming approximately $100J$ of energy with the target walking velocity as $0.1m/s$. Based on this, we conduct a linear search over energy thresholds of $20J$, $40J$, $60J$, $80J$, and $100J$ to identify the minimum threshold that enables the policy to achieve similar normalized episode length within 1500 iterations, as shown in Fig. 7. The $60J$ threshold provides the best trade-off between energy efficiency and locomotion stability. This approach is effective as it leverages physically meaningful thresholds while supporting parallel evaluation. A finer-grained search (e.g., with $10J$ increments) can discover policies with even lower energy consumption (e.g., $50J$) that still maintain stability. If further energy reduction is required, smaller step sizes enable finer control at the cost of additional computation.

5) **Hyperparameter Sensitivity Analysis:** To demonstrate that the proposed ECO framework requires less human effort in hyperparameter tuning, we conduct a sensitivity analysis by training ECO on a $0.1m/s$ linear velocity tracking task with




Fig. 7: **Energy constraint threshold analysis.** (a)(b): The mean energy cost and normalized episode length for different energy thresholds. Results averaged over 10 random seeds.

an energy threshold of $b_1 = 60J$. We fix the initial value of the Lagrange multiplier at 0 and vary the learning rate ($\beta_1 = 0.001, 0.01, 0.1$). As shown in Fig. 8(a), which illustrates the mean energy cost over iterations, all three learning rates converge to a similar energy cost of approximately $60J$ after an initial period of fluctuation. This pattern suggests that our Lagrangian approach ensures effective energy convergence across different learning rates. Fig. 8(b) shows the evolution of the Lagrange multiplier over iterations for different learning rates. A smaller learning rate causes the Lagrange multiplier to change more slowly, resulting in higher energy consumption initially due to constraint violations. As energy increases, the multiplier gradually rises, enabling faster convergence to the feasible region. Finally, Fig. 8(c) illustrates that our method achieves stable convergence across the range of learning rates.

6) **More Constraints Setting:** In previous work on quadruped robots [15, 47], additional reward terms such as self-collision, foot contact velocity, and foot clearance were integrated into the constraints to reduce the need for extensive reward scaling. We also attempted to incorporate these three constraints. Specifically, we trained an additional multi-head cost critic to estimate the cost returns for self-collision, foot contact velocity, and foot clearance. For these additional constraints, we used the converged reward values obtained from ECO (with two constraints) as the constraint thresholds. We analyze the performance of several algorithms under five




Fig. 8: **Hyperparameter sensitivity analysis.** The mean energy cost (a), Lagrange multiplier λ_1 (b), and normalized episode length (c) of ECO with the energy constraint $b_1 = 60J$, evaluated across different learning rates for λ_1 ($\beta_1 = 0.001, 0.01, 0.1$) during training. Results averaged over 10 random seeds.

constraints and investigate how the convergence properties of ECO are affected by different constraint configurations.

For the setup with five constraints, ECO is trained with a fixed initial Lagrange multiplier $\lambda_i = 0.0$ and a learning rate $\beta_i = 1e - 3$, while IPO and P3O used the same $\kappa = 1.0$ for each constraint. As shown in Fig. 9(a), we observed that the episode lengths for ECO, IPO, and P3O with five constraints are significantly reduced compared to ECO with only two constraints. Fig. 9(e) shows that P3O and IPO fail to converge on the foot velocity contact constraint, while ECO demonstrates better constraint satisfaction. Fig. 9(f) shows that all algorithms fail to converge on the self-collision cost. This is possibly due to predicting self-collision from the current privileged observation space is challenging.

In the four-constraint setting of ECO (without self-collision), we observe moderately better satisfaction of the energy constraint compared to the five-constraint case. Further removing the foot contact velocity constraint (three constraints setting) led to even longer episode lengths than the four-constraint setting, although performance remained below that of the original two-constraint configuration. Given the empirical results and our previous observation of the one-constraint setting—where omitting the symmetric reference motion constraint led to unnatural walking behaviors, as demonstrated in the accompanying video—we conclude that the two-constraint configuration (energy consumption and symmetry constraints) is the most suitable for achieving energy-efficient and consistently stable humanoid walking in our task. Further discussion and rationale are provided in Sec. VII-A.

TABLE V: **Algorithm performance comparison in the real world.** Aggregated over 10 runs, each lasting 10 seconds.

	Speed (m/s)	Energy (w)	Body Height (m)
MPC	0.124 ± 0.0072	9.1176 ± 5.0029	0.3483 ± 0.0056
PPO	0.126 ± 0.0088	3.9080 ± 1.8193	0.3745 ± 0.0034
ECO	0.122 ± 0.0089	1.7490 ± 0.8474	0.3790 ± 0.0030

VI. EXPERIMENTS

A. Real World Results

To verify the improvement in energy efficiency, we deploy the learned policy from ECO, along with the two baseline methods, on the real robot hardware BRUCE for comparison. Statistical results across multiple experiment runs, including the linear velocity along the x direction (speed), the total motor power (energy), and the body height in the world frame, are reported in Tab. V. Due to the sim-to-real gap, velocity tracking on the hardware is less accurate than in simulation. To ensure a fair comparison of energy consumption, we select three sets of data with similar, though not identical, speeds from multiple experiments using different methods. As shown in Fig. 10(a), ECO consistently maintains motor power near or below the energy cost threshold $2.5W(60J/24s)$ —approximately 6 times lower than MPC and 2.3 times lower than PPO according to the cumulative energy graph. Moreover, with the learned policy from ECO, BRUCE successfully rejects disturbances and traverses different outdoor terrains as illustrated in Fig. 11. The full hardware demonstration clips including comparisons with baselines, disturbance rejection, and outdoor walking can be seen in the accompanying video.

B. Emergent Behaviors

Empirical studies indicate that straight knee and heel-to-toe walking enhance energy efficiency in humans and humanoid robots [23, 62]. Instead of manually designing these behaviors for efficient locomotion, ECO shows that minimizing energy consumption can generate similar behaviors without explicit prescriptions, which also demonstrates that these bio-inspired behaviors are related to energy efficiency.

Fig. 12 illustrates the joint-level behaviors during walking. Fig. 12(a) compares the left knee joint angle over time, where ECO demonstrates a less flexed knee compared to MPC and PPO. In Fig. 12(b), the torque applied to the left knee joint is compared. ECO consistently requires lower torque to maintain stability and velocity tracking, resulting in reduced energy expenditure during walking. Due to the lack of precise force-torque sensors under the feet of BRUCE, the knee torques can also implicitly serve as a measure of the intensity of ground contact as the knee exerts the most force in humanoid locomotion. Additionally, we compare the body oscillation over a 10s stable walking period. Fig. 10(b) shows that BRUCE exhibits fewer oscillations with a lower frequency and smaller magnitude when walking with the policy from ECO, compared to MPC and PPO. The less body oscillation and smaller knee torques suggest that improved energy efficiency also results in lighter steps and reduced stomping.




Fig. 9: **Training performance with more constraints settings.** Comparison of training metrics for ECO, P3O, and IPO. The thresholds for energy consumption, mirror reference motion, foot clearance, foot contact velocity, and self-collision are set at $60J$, 0.05 , 3.6 , -528 , and 0.048 , respectively, as indicated by the **black dashed lines** in (b)-(f). Results averaged over 10 random seeds.




Fig. 10: **Motor energy consumption and body height oscillation in the real world.** Single-run deployment curves are shown for visual clarity. The relative body height is defined as the offset from the setpoint of the body height. The inset plot in (a) is the cumulative motor energy.

However, a heel-to-toe transition is not presented in the resulting policy as shown in Fig. 11. We argue that the heel-to-toe strategy, given the current hardware design of BRUCE, demands additional control effort to execute effectively. Furthermore, it requires a higher swing leg lift, which increases energy consumption in the hip and knee motors.




Fig. 11: **Screenshots of the sim-to-real transfer.** BRUCE successfully rejects disturbance and traverses outdoor terrains which can be seen in the accompanying video.

VII. DISCUSSION

A. Rewards v.s. Constraints

In this study, we find that tuning reward scales through hyperparameter search within the RL framework (PPO) is an inefficient approach to reducing energy consumption in humanoid robots. This inefficiency arises because reward scales do not correspond linearly to physical quantities, demanding a delicate balance between meeting energy requirements, following task commands, and maintaining motion stability. In contrast, constrained RL explores a space defined by explicit constraints, many of which have direct physical significance, such as energy consumption limits and joint restrictions. This makes tuning constraints more straightforward and effective compared to adjusting reward or loss weights.

In this work, we try to incorporate self-collision, foot contact velocity, and foot clearance as additional constraints. Although such constraints have been successfully applied to quadruped robots [15, 47], we find it challenging to transfer these approaches directly to humanoids. This difficulty arises not only from technical challenges associated with the




Fig. 12: **Comparison of left knee joint angle and torque between different methods in the real world.** We present single-run deployment curves for visual clarity. A larger knee joint angle reflects a more flexed knee during walking, while a smaller knee joint torque at ground contact indicates a lighter step.

heterogeneity of constraint signals but also from differences in morphology and balance dynamics. First, the cost terms associated with different constraints often differ in magnitude, which can lead to imbalanced gradient updates when training the multi-head cost critic [15]. As a result, costs with smaller magnitudes—such as the self-collision cost in Fig. 9(f)—tend to converge more slowly. This issue could be mitigated by applying appropriate normalization techniques [63] to rescale cost signals to comparable magnitudes. Second, accurately predicting different types of costs may require heterogeneous information. For instance, inferring self-collision from the current privileged observation (i.e., joint angles and base-link kinematics) is particularly challenging, as it requires the network to implicitly reconstruct the geometry of the body in the world frame from joint-centric data. This suggests that a richer observation space—including world-frame information—or a specialized cost critic for self-collision may be necessary to improve both prediction accuracy and constraint satisfaction.

Furthermore, humanoid robots have much smaller support polygons than quadruped robots, especially in the single-support phase, where the support region further shrinks. This renders the space of feasible trajectories much smaller than quadruped robots. In this case, imposing a large amount of feasibility constraints in constrained RL may complex the search space. We believe further study is required to handle a larger number of constraints in humanoid locomotion. For instance, incorporating model-based prior knowledge to refine the search space—such as using nullspace projections to prioritize control tasks, as commonly done in the control community [64]—could help in selecting appropriate constraints and mitigating task conflicts. A comprehensive investigation

of such approaches is beyond the scope of this work and is left for future research.

B. Optimization Method Selection

Our experiments show that among the four constrained RL algorithms, PPO-Lag is the most effective. CRPO, as a primal algorithm, does not introduce additional dual variables for optimization and involves fewer hyperparameters, making it theoretically simpler and easier to implement compared to primal-dual methods. However, it lacks the smooth transition between reward and constraint optimization objectives that primal-dual methods offer, making it difficult to achieve stable convergence in humanoid robots. P3O uses clamping functions on constraints to apply a linear cost gradient penalty until the constraints are met. However, this linear penalty makes balancing multiple constraints difficult, as each constraint has distinct physical significance and magnitude, requiring separate coefficients κ_i^{P3O} . Small coefficients lead to slow convergence, while large ones make the system overly conservative. IPO, as a method that theoretically also employs penalty functions, exhibits performance similar to P3O in our experiments. It also requires careful tuning of the coefficient κ_i^{IPO} to balance constraints effectively.

In contrast, ECO leverages PPO-Lag to update both the Lagrange multipliers and the policy, allowing the Lagrange multipliers and their learning rates to adapt to the magnitude of each constraint. By utilizing the Adam optimizer to update these multipliers and learning rates, ECO effectively adapts to each constraint while smoothly transitioning the optimization objective between rewards and constraints. As a result, ECO achieves stable convergence in humanoid robots.

However, ECO faces challenges in optimizing the five constraints considered in our study, highlighting the need for improved optimization techniques in future work. Leveraging the dynamics model of humanoid robots to improve the convergence of constrained RL within a restricted feasible region could be a promising direction [13, 65, 66].

C. Sim-to-Real Failure Cases

Our methods successfully transfer to the real world, and we identify several key factors that critically affect policy transfer:

- 1) **PD Gain Tuning:** We find that carefully tuning the proportional (K_p) and derivative (K_d) gains of the low-level PD controller is essential to align the joint-position tracking performance between simulation and hardware. Otherwise, differences in system dynamics can cause the policy to produce divergent behaviors in the real world.
- 2) **Compliant ankle strategy:** We find that high PD gains at the ankle lead to instability upon ground contact due to sensor noise. To address this, we reduce the ankle’s PD gains, making the joints more compliant and the policy less sensitive to noise. This simple adjustment significantly enhances robustness to terrain variations and modeling errors in real-world tests.
- 3) **Sensor Noise and Bias:** Real sensors exhibit both high-frequency noise (especially in joint-velocity measurements) and low-frequency biases (e.g. Euler-angle drift

upon reset). To account for this, we inject noise and random biases—matched in magnitude to hardware measurements—into the corresponding state observations during domain randomization.

D. More complex settings

While our method is effective for stable walking, applying it to more complex scenarios, such as transitioning from flat surfaces to stairs, presents new challenges. For instance, a lower foot lift is energy-efficient on flat surfaces, but stair climbing requires higher foot lifts and precise planning to navigate varying stair heights efficiently. One possible extension is to introduce mode-aware constraints [67, 68] that adapt according to different terrains and command types. For example, energy constraints could be relaxed when stair climbing is anticipated. A terrain classifier or visual encoder [69–71] could be used to identify terrain types, allowing the policy to dynamically adjust constraint thresholds accordingly. In addition, model-based planning [13, 72, 73] can be incorporated, in which a high-level planner selects locomotion modes based on the environment, while the low-level control executes constrained locomotion policy conditioned on the selected mode. This hierarchical approach aligns well with the constrained RL framework, as constraints can be defined and enforced differently across modes. We consider this an exciting future direction and plan to explore model-based extensions in our follow-up work.

VIII. CONCLUSION

In this paper, we optimize the energy consumption for humanoid walking through constrained RL by separating energy consumption and reference motion objectives from the reward function. We demonstrate that the ECO framework can converge to a policy that achieves lower energy consumption while maintaining robustness. Meanwhile, ECO is resilient to hyperparameter settings, thereby eliminating the need for extensive tuning on the weights of rewards. Extensive experiments in both simulation and real-world hardware validate that our approach can significantly improve the energy efficiency of humanoid walking without compromising robustness.


REFERENCES

- [1] Z. Wang, Y. Jia, L. Shi, H. Wang, H. Zhao, X. Li, J. Zhou, J. Ma, and G. Zhou, “Arm-constrained curriculum learning for loco-manipulation of the wheel-legged robot,” *arXiv preprint arXiv:2403.16535*, 2024.
- [2] Z. He, J. Wu, J. Zhang, S. Zhang, Y. Shi, H. Liu, L. Sun, Y. Su, and X. Leng, “Cdm-mpc: An integrated dynamic planning and control framework for bipedal robots jumping,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 9, no. 7, pp. 6672–6679, 2024.
- [3] J. Zhang, J. Shen, Y. Liu, and D. Hong, “Design of a jumping control framework with heuristic landing for bipedal robots,” in *IEEE/RAS International Conference on Intelligent Robots and Systems (IROS)*, pp. 8502–8509, IEEE, 2023.
- [4] H. Liu, Q. Xie, Z. Zhang, T. Yuan, S. Wang, Z. Wang, X. Leng, L. Sun, J. Zhang, Z. He, *et al.*, “Pr2: A physics-and photo-realistic humanoid testbed with pilot study in competition,” *Journal of Field Robotics*, 2025.
- [5] P.-B. Wieber, R. Tedrake, and S. Kuindersma, “Modeling and control of legged robots,” in *Springer handbook of robotics*, pp. 1203–1234, Springer, 2016.
- [6] Y. G. Alqaham, J. Cheng, and Z. Gan, “Energy-optimal asymmetrical gait selection for quadrupedal robots,” *IEEE Robotics and Automation Letters (RA-L)*, 2024.
- [7] Z. Fu, A. Kumar, J. Malik, and D. Pathak, “Minimizing energy consumption leads to the emergence of gaits in legged robots,” *arXiv preprint arXiv:2111.01674*, 2021.
- [8] Z. Fu, Q. Zhao, Q. Wu, G. Wetzstein, and C. Finn, “Humanplus: Humanoid shadowing and imitation from humans,” *arXiv preprint arXiv:2406.10454*, 2024.
- [9] X. Gu, Y.-J. Wang, and J. Chen, “Humanoid-gym: Reinforcement learning for humanoid robot with zero-shot sim2real transfer,” *arXiv preprint arXiv:2404.05695*, 2024.
- [10] S. H. Jeon, S. Heim, C. Khazoom, and S. Kim, “Benchmarking potential based rewards for learning humanoid locomotion,” in *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 9204–9210, IEEE, 2023.
- [11] Y. Chow, M. Ghavamzadeh, L. Janson, and M. Pavone, “Risk-constrained reinforcement learning with percentile risk criteria,” *Journal of Machine Learning Research*, vol. 18, no. 167, pp. 1–51, 2018.
- [12] T. He, W. Zhao, and C. Liu, “Autocost: Evolving intrinsic cost for zero-violation reinforcement learning,” in *AAAI Conference on Artificial Intelligence*, vol. 37, pp. 14847–14855, 2023.
- [13] W. Huang, J. Ji, B. Zhang, C. Xia, and Y. Yang, “Safe dreamerv3: Safe reinforcement learning with world models,” *arXiv preprint arXiv:2307.07176*, 2023.
- [14] J. Achiam, D. Held, A. Tamar, and P. Abbeel, “Constrained policy optimization,” in *International conference on machine learning*, pp. 22–31, PMLR, 2017.
- [15] Y. Kim, H. Oh, J. Lee, J. Choi, G. Ji, M. Jung, D. Youm, and J. Hwangbo, “Not only rewards but also constraints: Applications on legged robot locomotion,” *IEEE Transactions on Robotics (T-RO)*, 2024.
- [16] S. Collins, A. Ruina, R. Tedrake, and M. Wisse, “Efficient bipedal robots based on passive-dynamic walkers,” *Science*, vol. 307, no. 5712, pp. 1082–1085, 2005.
- [17] R. Tedrake, T. W. Zhang, M.-f. Fong, and H. S. Seung, “Actuating a simple 3d passive dynamic walker,” in *IEEE International Conference on Robotics and Automation (ICRA)*, vol. 5, pp. 4656–4661, IEEE, 2004.
- [18] S. Kuindersma, R. Deits, M. Fallon, A. Valenzuela, H. Dai, F. Permenter, T. Koolen, P. Marion, and R. Tedrake, “Optimization-based locomotion planning, estimation, and control design for the atlas humanoid robot,” *Autonomous robots*, vol. 40, pp. 429–455, 2016.
- [19] J. Shen, J. Zhang, Y. Liu, and D. Hong, “Implementation of a robust dynamic walking controller on a miniature bipedal robot with proprioceptive actuation,” in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, pp. 39–46, IEEE, 2022.
- [20] W. Xi, Y. Yesilevskiy, and C. D. Remy, “Selecting gaits for economical locomotion of legged robots,” *International Journal of Robotics Research (IJRR)*, vol. 35, no. 9, pp. 1140–1154, 2016.
- [21] B. Ponton, M. Khadiv, A. Meduri, and L. Righetti, “Efficient multi-contact pattern generation with sequential convex approximations of the centroidal dynamics,” *IEEE Transactions on Robotics (T-RO)*, vol. 37, no. 5, pp. 1661–1679, 2021.
- [22] J. Wang, S. Kim, T. S. Lembono, W. Du, J. Shim, S. Samadi, K. Wang, V. Ivan, S. Calinon, S. Vijayakumar, *et al.*, “Online multi-contact receding horizon planning via value function approximation,” *IEEE Transactions on Robotics (T-RO)*, vol. 40, pp. 2791–2810, 2024.
- [23] S. Fasano, J. Foster, S. Bertrand, C. DeBuys, and R. Griffin, “Efficient, dynamic locomotion through step placement with straight legs and rolling contacts,” in *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1143–1150, IEEE, 2024.
- [24] W. Cui, S. Li, H. Huang, B. Qin, T. Zhang, L. Zheng, Z. Tang, C. Hu, N. Yan, J. Chen, *et al.*, “Adapting humanoid locomotion over challenging terrain via two-phase training,” in *8th Annual Conference on Robot Learning*.
- [25] K. Li, G. Lai, and X. Yao, “Interactive evolutionary multiobjective optimization via learning to rank,” *IEEE Transactions on Evolutionary Computation*, vol. 27, no. 4, pp. 749–763, 2023.
- [26] I. Radosavovic, B. Zhang, B. Shi, J. Rajasegaran, S. Kamat, T. Darrell, K. Sreenath, and J. Malik, “Humanoid locomotion as next token prediction,” in *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [27] J. Siekmann, Y. Godse, A. Fern, and J. Hurst, “Sim-to-real learning of all common bipedal gaits via periodic reward composition,” in *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 7309–7315, IEEE, 2021.
- [28] A. Kumar, Z. Li, J. Zeng, D. Pathak, K. Sreenath, and J. Malik, “Adapting rapid motor adaptation for bipedal robots,” in *IEEE/RAS International Conference on Intelligent Robots and Systems (IROS)*, pp. 1161–1168, IEEE, 2022.


- [29] H. J. Lee, S. Hong, and S. Kim, "Integrating model-based footstep planning with model-free reinforcement learning for dynamic legged locomotion," *arXiv preprint arXiv:2408.02662*, 2024.
- [30] S. H. Bang, C. A. Jove, and L. Sentis, "RL-augmented mpc framework for agile and robust bipedal footstep locomotion planning and control," *arXiv preprint arXiv:2407.17683*, 2024.
- [31] W. Yu, G. Turk, and C. K. Liu, "Learning symmetric and low-energy locomotion," *ACM Transactions on Graphics (TOG)*, vol. 37, no. 4, pp. 1–12, 2018.
- [32] Z. Su, X. Huang, D. Ordonez-Apaez, Y. Li, Z. Li, Q. Liao, G. Turrisi, M. Pontil, C. Semini, Y. Wu, and K. Sreenath, "Leveraging symmetry in rl-based legged locomotion control," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 6899–6906, 2024.
- [33] 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.
- [34] 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.
- [35] J. Xu, K. Li, and M. Abusara, "Preference based multi-objective reinforcement learning for multi-microgrid system optimization problem in smart grid," *Mementic Computing*, vol. 14, no. 2, pp. 225–235, 2022.
- [36] K. Li and H. Guo, "Human-in-the-loop policy optimization for preference-based multi-objective reinforcement learning," *arXiv preprint arXiv:2401.02160*, 2024.
- [37] E. Altman, *Constrained Markov decision processes: stochastic modeling*. Routledge, 1999.
- [38] S. Li, K. Li, W. Li, and M. Yang, "Evolutionary alternating direction method of multipliers for constrained multi-objective optimization with unknown constraints," *IEEE Transactions on Evolutionary Computation*, 2024.
- [39] S. Wang and K. Li, "Constrained bayesian optimization under partial observations: Balanced improvements and provable convergence," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, pp. 15607–15615, 2024.
- [40] S. Paternain, L. F. O. Chamon, M. Calvo-Fullana, and A. Ribeiro, "Constrained reinforcement learning has zero duality gap," *CoRR*, vol. abs/1910.13393, 2019.
- [41] A. Ray, J. Achiam, and D. Amodei, "Benchmarking safe exploration in deep reinforcement learning," *arXiv preprint arXiv:1910.01708*, vol. 7, no. 1, p. 2, 2019.
- [42] P. Liu, D. Tateo, H. B. Ammar, and J. Peters, "Robot reinforcement learning on the constraint manifold," in *Conference on Robot Learning*, pp. 1357–1366, PMLR, 2022.
- [43] P. Kicki, P. Liu, D. Tateo, H. Bou-Ammar, K. Walas, P. Skrzypczynski, and J. Peters, "Fast kinodynamic planning on the constraint manifold with deep neural networks," *IEEE Transactions on Robotics (T-RO)*, vol. 40, pp. 277–297, 2023.
- [44] T. Xu, Y. Liang, and G. Lan, "Crpo: A new approach for safe reinforcement learning with convergence guarantee," in *International Conference on Machine Learning*, pp. 11480–11491, PMLR, 2021.
- [45] Y. Liu, J. Ding, and X. Liu, "Ipo: Interior-point policy optimization under constraints," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 34, pp. 4940–4947, 2020.
- [46] L. Zhang, L. Shen, L. Yang, S. Chen, B. Yuan, X. Wang, and D. Tao, "Penalized proximal policy optimization for safe reinforcement learning," *arXiv preprint arXiv:2205.11814*, 2022.
- [47] J. Lee, L. Schroth, V. Klemm, M. Bjelonic, A. Reske, and M. Hutter, "Exploring constrained reinforcement learning algorithms for quadrupedal locomotion," in *IEEE/RAS International Conference on Intelligent Robots and Systems (IROS)*, pp. 11132–11138, IEEE, 2024.
- [48] J. Ji, B. Zhang, J. Zhou, X. Pan, W. Huang, R. Sun, Y. Geng, Y. Zhong, J. Dai, and Y. Yang, "Safety gymnasium: A unified safe reinforcement learning benchmark," *Advances in Neural Information Processing Systems*, vol. 36, 2023.
- [49] X. Gu, Y.-J. Wang, X. Zhu, C. Shi, Y. Guo, Y. Liu, and J. Chen, "Advancing humanoid locomotion: Mastering challenging terrains with denoising world model learning," *arXiv preprint arXiv:2408.14472*, 2024.
- [50] D. P. Kingma, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [51] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [52] J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel, "High-dimensional continuous control using generalized advantage estimation," *arXiv preprint arXiv:1506.02438*, 2015.
- [53] 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*, pp. 91–100, PMLR, 2022.
- [54] R. J. Williams, "Simple statistical gradient-following algorithms for connectionist reinforcement learning," *Machine learning*, vol. 8, pp. 229–256, 1992.
- [55] Y. Liu, J. Shen, J. Zhang, X. Zhang, T. Zhu, and D. Hong, "Design and control of a miniature bipedal robot with proprioceptive actuation for dynamic behaviors," in *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 8547–8553, IEEE, 2022.
- [56] BRUCE, A kid-size humanoid robot open-platform for research and education. <https://www.westwoodrobotics.io/bruce/>. Accessed: 2024-08-30.
- [57] J. Liang, V. Makoviychuk, A. Handa, N. Chentanez, M. Macklin, and D. Fox, "Gpu-accelerated robotic simulation for distributed reinforcement learning," in *Conference on Robot Learning*, PMLR, 2018.
- [58] L. Pinto, M. Andrychowicz, P. Welinder, W. Zaremba, and P. Abbeel, "Asymmetric actor critic for image-based robot learning," *arXiv preprint arXiv:1710.06542*, 2017.
- [59] R. P. Singh, Z. Xie, P. Gergondet, and F. Kanehiro, "Learning bipedal walking for humanoids with current feedback," *IEEE Access*, 2023.
- [60] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *IEEE/RAS International Conference on Intelligent Robots and Systems (IROS)*, pp. 5026–5033, IEEE, 2012.
- [61] N. Koenig and A. Howard, "Design and use paradigms for gazebo, an open-source multi-robot simulator," in *IEEE/RAS International Conference on Intelligent Robots and Systems (IROS)*, vol. 3, pp. 2149–2154, IEEE, 2004.
- [62] P. G. Adamczyk, S. H. Collins, and A. D. Kuo, "The advantages of a rolling foot in human walking," *Journal of experimental biology*, vol. 209, no. 20, pp. 3953–3963, 2006.
- [63] D. Hafner, J. Pasukonis, J. Ba, and T. Lillicrap, "Mastering diverse domains through world models," *arXiv preprint arXiv:2301.04104*, 2023.
- [64] A. Dietrich, C. Ott, and A. Albu-Schaffer, "An overview of null space projections for redundant, torque-controlled robots," *International Journal of Robotics Research (IJRR)*, vol. 34, no. 11, pp. 1385–1400, 2015.
- [65] J. Levy, T. Westenbroek, and D. Fridovich-Keil, "Learning to walk from three minutes of real-world data with semi-structured dynamics models," *arXiv preprint arXiv:2410.09163*, 2024.
- [66] Y. As, I. Usmanova, S. Curi, and A. Krause, "Constrained policy optimization via bayesian world models," *arXiv preprint arXiv:2201.09802*, 2022.
- [67] J.-P. Sleiman, F. Farshidian, and M. Hutter, "Versatile multicontact planning and control for legged loco-manipulation," *Science Robotics*, vol. 8, no. 81, p. eadg5014, 2023.
- [68] M. Zhang, Y. Ma, T. Miki, and M. Hutter, "Learning to open and traverse doors with a legged manipulator," *arXiv preprint arXiv:2409.04882*, 2024.
- [69] A. Agarwal, A. Kumar, J. Malik, and D. Pathak, "Legged locomotion in challenging terrains using egocentric vision," in *Conference on robot learning*, pp. 403–415, PMLR, 2023.
- [70] E. Chane-Sane, J. Amigo, T. Flayols, L. Righetti, and N. Mansard, "Soloparkour: Constrained reinforcement learning for visual locomotion from privileged experience," in *Conference on Robot Learning*, arXiv, 2024.
- [71] W. Yu, D. Jain, A. Escontrela, A. Iscen, P. Xu, E. Coumans, S. Ha, J. Tan, and T. Zhang, "Visual-locomotion: Learning to walk on complex terrains with vision," in *5th Annual Conference on Robot Learning*, 2021.
- [72] N. Hansen, J. SV, V. Sobal, Y. LeCun, X. Wang, and H. Su, "Hierarchical world models as visual whole-body humanoid controllers," *arXiv preprint arXiv:2405.18418*, 2024.
- [73] Y. Ma, F. Farshidian, T. Miki, J. Lee, and M. Hutter, "Combining learning-based locomotion policy with model-based manipulation for legged mobile manipulators," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2377–2384, 2022.



Weidong Huang received the M.S. degree from the Beihang University in 2024. He is now a research engineer at State Key Laboratory of General Artificial Intelligence, Beijing Institute for General Artificial Intelligence (BIGAI). He received his B.S. degree in South China Normal University in 2021. His research interests include robotics, reinforcement learning, and model-based planning.



Jingwen Zhang (Member, IEEE) received the B.S. degree and the M.S. degree from the Department of Mechanical Engineering of Tsinghua University and Columbia University in 2015 and 2017, and the Ph.D. degree from the Department of Mechanical and Aerospace Engineering, University of California, Los Angeles in 2023. He is now a research scientist at State Key Laboratory of General Artificial Intelligence, Beijing Institute for General Artificial Intelligence (BIGAI). His research interests include robotics, optimal control, and motion planning.



Jiongye Li is now a third-year student in the Department of Automation of Tsinghua University. He is now an intern at State Key Laboratory of General Artificial Intelligence, Beijing Institute for General Artificial Intelligence (BIGAI). His research interests include humanoid robots, reinforcement learning and optimal control.



Shibowen Zhang received the B.E. degree in Control Science and Engineering, from Tongji University, Shanghai, China, in 2024. He is currently a first-year Ph.D. student at the School of Control Science and Engineering, University of Science and Technology of China (USTC). His research interests include robotics and control.



Jiayang Wu is currently a senior undergraduate student in the Department of Computer Science and Technology at Harbin Institute of Technology (HIT). He is also an intern research assistant at the Beijing Institute for General Artificial Intelligence (BIGAI). His research interests include robotics, optimal control, and reinforcement learning.



Jiayi Wang (Member, IEEE) received his Ph.D. degree in robotics from the University of Edinburgh in 2023. He is currently a Research Scientist at Beijing Institute for General Artificial Intelligence (BIGAI). Before that, he was a research associate at the University of Edinburgh in 2024 and was a research intern at the Italian Institute of Technology in 2017. His research interests include humanoid locomotion, multi-contact motion planning, reinforcement learning, and trajectory optimization.



and cognitive robotics.

Hangxin Liu (Member, IEEE) received his Ph.D. degree in Computer Science from the University of California, Los Angeles (UCLA) in 2021. He is currently a research scientist at State Key Laboratory of General Artificial Intelligence, Beijing Institute for General Artificial Intelligence (BIGAI). He received an M.S. degree in Mechanical Engineering from UCLA in 2018 and two B.S. degrees in Mechanical Engineering and Computer Science, both from Virginia Tech in 2016. His research interests focus on robot perception, learning, human-robot interaction,



artificial general collective intelligence through multi-agent reinforcement learning. He has maintained a track record of more than sixty publications at top conferences (NeurIPS, ICML, ICLR, etc) and top journals (Artificial Intelligence, National Science Review, etc), along with the best system paper award at CoRL 2020 and the best blue-sky paper award at AAMAS 2021. He was awarded ACM SIGAI China Rising Star and World AI Conference (WAIC'22) Rising Star.

Yaodong Yang (Member, IEEE) received the bachelor's degree from the University of Science and Technology of China, the MSc degree from Imperial College London, and the PhD degree from University College London (nominated by UCL for Joint AAAI/ACM SIGAI Doctoral Dissertation Award). He is an assistant professor with Institute for AI, Peking University. Before joining Peking University, he was an assistant professor with King's College London. He studies game theory, reinforcement learning and multi-agent systems, aiming to achieve



Yao Su (Member, IEEE) received the B.S. degree from the School of Mechatronic Engineering, Harbin Institute of Technology in 2016, and the M.S. and Ph.D. degrees from the Department of Mechanical and Aerospace Engineering, University of California, Los Angeles (UCLA) in 2017 and 2021. He is now a research scientist at State Key Laboratory of General Artificial Intelligence, Beijing Institute for General Artificial Intelligence (BIGAI). His research interests include robotics, control, optimization, trajectory planning, and mechatronics.