Article | Received 29 March 2026; Revised 14 May 2026; Accepted 8 June 2026; Published 27 July 2026
https://doi.org/10.55092/rl20260023

Impact-aware compliant foot placement for quiet humanoid locomotion via reinforcement learning

Kai Jian[1], Shaojie Zhang[1], Dong Zhu[2] and Xiangyu Shao[1,*]

[1] State Key Laboratory of Robotics and System (HIT), and School of Astronautics, Harbin Institute of Technology, Harbin, China

[2] Chengdu CHG Robots & Intelligent Equipment Institute Co., Ltd., Chengdu, China

* Correspondence author; E-mail: xiangyushao@hit.edu.cn.

Highlights:

- A virtual compliance parameter modulation strategy is proposed to achieve compliant landing control for robots, enabling adjustable quiet walking performance across diverse indoor environments.
- A foot corner contact modeling method is developed to effectively prevent premature unilateral contact, thereby ensuring smoother and more stable foot touchdown during locomotion.
- The effectiveness of the proposed approach is validated through both simulation and real-robot experiments, demonstrating its robustness and practical applicability in real indoor scenarios.

Abstract: Recent advances in legged robots have substantially improved their locomotion capabilities over outdoor and complex terrains. However, quiet locomotion for humanoid robots in noise-sensitive indoor environments remains underexplored, despite its growing importance in human-centered applications. While encouraging progress has been made in quadrupedal robots, transferring the quiet locomotion ability to humanoid robots remains nontrivial due to their fundamentally different foot-ground contact patterns. This paper proposes a control method for reducing foot–ground contact noise during humanoid walking, achieving compliant contact and continuous regulation of locomotion noise by establishing a foot corner contact model along with a virtual compliance parameter. The experimental results show that the average sound pressure level is reduced by 4.88 dB.

Keywords: humanoid robots; reinforcement learning; quiet locomotion; foot contact modeling

1. Introduction

Legged robots have garnered significant attention in recent years, particularly due to their ability to achieve more flexible and natural motion in complex and unstructured terrains through reinforcement learning–based motion control [1–5]. This adaptability has enabled robots to stably and autonomously perform practical tasks such as search and rescue [6], inspection [7], and service industries [8].

Despite remarkable achievements in outdoor and industrial settings [9–12], the application of legged robots, particularly humanoid robots, in noise-sensitive indoor environments has not been extensively investigated. Indoor environments such as hospitals, elderly care facilities, offices, and homes demand low-noise operation and intuitive interaction. In these indoor environments, the continuous mechanical noise generated by robot movement can be disruptive, thus limiting the practical deployment of legged robots in these environments.

Motivated by these concerns, the robotics community has begun to explore sound engineering as an emerging research direction [13]. Existing studies primarily focus on robotic manipulators [14,15], quadruped robots [16,17], and wheeled robots [18]. These works research how robot-generated sounds influence human comfort, perception, and trust, as well as propose strategies to reduce mechanical noise. For instance, Watanabe *et al.* [16] present a reinforcement learning framework in which the policy outputs not only target joint position but also proportional–derivative (PD) gain scales, enabling dynamic adaptation of low-level joint control parameters to suppress locomotion noise in quadruped robots. Similarly, Cao *et al.* [17] introduce the Minimizing acoUsTic noisE (MUTE) framework, which incorporates a silence factor to explicitly regulate the trade-off between quiet locomotion and velocity tracking performance. These approaches are effective for robots with relatively small contact areas but generally rely on simplified point-contact assumptions that do not explicitly model distributed impact behavior during touchdown.

During locomotion, humanoid robots generate noise from multiple sources. Previous studies on acoustic locomotion in legged robots have reported that locomotion noise mainly arises from repetitive foot-ground impacts and the associated structural vibration during walking [17,19]. In comparison, actuator and transmission noise are generally less significant under typical indoor locomotion conditions [17]. Since impact-induced acoustic energy is closely related to touchdown velocity and contact force [19], the present work primarily focuses on reducing impact-related locomotion noise through touchdown dynamics regulation. Other acoustic sources are not explicitly isolated in this study and remain part of the overall measured locomotion noise. Therefore, the problem of quiet locomotion can be formulated as enforcing more compliant robot–ground interactions. This perspective is consistent with prior work on soft landing control [20], where impact forces are mitigated through trajectory optimization and learning-based strategies. The two quadruped examples [16,17] mentioned earlier follow a similar idea. Additionally, compliant control can mitigate impact forces, thereby contributing to the protection of the robot.

However, previous works have not thoroughly explored the acoustic impact of humanoid robot footsteps. The regulation of compliant contact in humanoid robots remains insufficiently investigated [21]. A common strategy for humanoid robot locomotion using ankle force-torque sensors is to generate compliant joint motions from the Ground Reaction Wrench (GRW) to distribute the contact forces along the sole [22–24]. However, these methods rely on feedback after contact with the ground to adjust the torque, and cannot adjust the speed before contact. Guadarrama-Olvera *et al.* [25] propose a method that can adjust the speed in advance, but it relies on a large number of tactile pixels attached to the foot. Introducing external sensors increases cost. Unlike existing methods, the proposed approach is based on rigid body kinematics and explicitly models the ground contact velocity of a single foot corner. It can adjust the speed and joint torque in advance without introducing external sensors, thereby improving local impact suppression and motion noise reduction in humanoid walking.

This work focuses on reducing noise by regulating impact-related dynamics through control strategies rather than modifying the mechanical structure. Unlike quadruped robots that primarily rely on localized point contacts, humanoid robots interact with the ground using large flat soles, resulting in distributed surface contact [26]. This fundamental difference alters the touchdown dynamics substantially. Quadruped locomotion typically concentrates impact forces near a single contact point, leading to relatively simple contact transitions and localized shock propagation. In contrast, humanoid robots often experience asynchronous and spatially distributed foot–ground contacts due to ankle rotation and posture adjustment, where the toe, heel, or lateral edges may contact the ground at different instants [27]. Consequently, impact forces are unevenly distributed across the foot surface, introducing coupled translational and rotational impact dynamics. Small variations in foot orientation or angular velocity can produce significant localized impact peaks, especially during heel-strike or toe-strike events. These distributed contact behaviors increase the difficulty of compliant contact regulation and make humanoid locomotion more sensitive to localized dynamics than quadruped systems. Therefore, regulating only the center-of-mass motion of the foot is insufficient for accurate impact characterization, which motivates the proposed foot corner contact modeling strategy.

Based on the above analysis, this work proposes a control approach for compliant impact mitigation in humanoid locomotion under fixed hardware configurations. Rather than directly minimizing noise as an isolated objective, the proposed method explicitly regulates pre-contact motion states to reduce impulsive forces during touchdown. By optimizing motion generation and contact modulation strategies, smooth and compliant walking behavior is achieved. A series of high-fidelity simulations and real-world experiments conducted in representative indoor environments validate the effectiveness of the approach. Experimental results demonstrate that the proposed method can substantially attenuate impact forces and significantly reduce both average and peak locomotion noise levels, while maintaining walking stability, performance robustness, and dynamic consistency. The main contributions of this work are summarized as follows:

(1) A noise-minimization method is proposed by introducing a virtual compliance parameter into leg motion control, which also mitigates impact forces during foot–ground contact. By adjusting this parameter, different levels of quiet locomotion can be achieved to accommodate diverse indoor application scenarios.

(2) A foot corner contact modeling strategy is adopted, which does not rely on the velocity of the foot center of mass. This approach effectively prevents premature unilateral foot contact and ensures smoother and more stable foot touchdown during walking.

(3) The effectiveness of the proposed method is validated through both simulation and real world experiments. Compared with baseline approaches, the proposed method achieves an average reduction of 4.88 dB in locomotion noise, demonstrating its effectiveness in achieving compliant contact in real indoor environments.

The rest of this paper is organized as follows: Section 2 elaborates on our proposed foot corner contact modeling strategy and the virtual compliance parameter motion control method based on reinforcement learning. Simulation and real world experiment results are given in Section 3, and Section 4 concludes this paper.

2. Methods

The amplitude of sound generated by an impact is typically related to the impact energy transferred during the contact process. For the foot–ground contact process, the kinetic energy of the foot can be expressed as $E = \frac{1}{2}mv^2$, indicating that touchdown velocity is an important factor influencing impact-induced excitation. However, the final perceived locomotion noise is also affected by multiple coupled factors, including foot–ground material properties, contact impedance, structural vibration transmission, robot frame resonance, and actuator-related acoustic effects. Consequently, the relationship between touchdown velocity and perceived acoustic amplitude is not strictly deterministic. While noise generated by motors and joint mechanisms is influenced by the robot's structural design and can be reduced through optimization techniques, such as minimizing mechanical friction [28], it is not the primary focus of this study. In this work, the objective is not to redesign the robot mechanical structure or optimize material-level acoustic properties, but rather to reduce impact-related excitation through motion control and touchdown dynamics regulation under a fixed hardware configuration. All experiments are conducted under identical conditions, including the same robot hardware, foot structure, floor surface, and environment, to ensure a fair comparison. Under these controlled conditions, the sound amplitude can be approximately associated with the foot velocity at the instant of contact.

Considering the complexity of explicit acoustic modeling in physical simulation, this study addresses the problem by imposing constraints on the foot velocity prior to ground contact. By suppressing the impact velocity at touchdown, the proposed approach indirectly regulates the noise generated during foot–ground collisions as well as the impact forces with the ground.

Since only proprioceptive sensing is available, the problem is formulated as a partially observable Markov decision process (POMDP), defined by the tuple $M = \{S, A, T, R, \Omega, O, \gamma\}$. At each time step, the agent selects an action $a_t \in A$ based on the current state $s_t \in S$, transitions to the next environment state according to the transition function $T(s_{t+1}|s_t, a_t)$, receives a reward r_t from the reward function $R(s_t, a_t)$, and obtains an observation $o_t \in \Omega$ through the observation function $O(o_t|s_{t+1}, a_t)$. The objective is to learn an optimal policy π^* that maximizes the expected cumulative reward $J(\pi)$, which can be expressed as follows:

$$J(\pi) = \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t) \right] \tag{1}$$

where γ is the discount factor.

2.1. Foot corner contact modeling

In humanoid locomotion, touchdown impact dynamics are influenced not only by the translational motion of the foot rigid body but also by rotational motion around the ankle joints. During the touchdown phase, ankle pitch rotation can produce significantly different local velocities across the foot surface [27]. Consequently, even when the center-of-mass velocity of the foot remains relatively small, the toe or heel regions may still experience large instantaneous downward velocities, resulting in concentrated local impact forces and abrupt contact transitions.

Under such conditions, representing foot touchdown behavior solely using the center-of-mass velocity is insufficient to accurately characterize localized impact dynamics. In particular, center-based representations cannot distinguish asymmetric contact events such as toe-first or heel-first landing patterns, which are common in humanoid walking and often generate sharp impact peaks and increased locomotion noise.

To more precisely characterize the local contact behavior between the foot and the ground, this work introduces a foot corner contact modeling approach. The proposed foot corner contact modeling framework is illustrated in Figure 1. Specifically, four virtual corner points are defined in the rigid-body coordinate frame of each foot, corresponding to the front-left, front-right, rear-left, and rear-right corners of the foot sole. Using the instantaneous linear and angular velocities of the foot, the velocities of these corner points in the world frame are computed through rigid-body kinematic relationships. The corresponding formulation is given as follows:

$$v_i = v_{com} + \omega \times r_i \tag{2}$$

In the above equation, v_i denotes the velocity of the i-th foot corner point, v_{com} represents the linear velocity of the foot rigid body's center of mass, ω denotes the angular velocity of the foot rigid body, and r_i is the position vector from the center of mass to the corresponding foot corner point. The second term $\omega \times r_i$ represents the velocity contribution induced by rotational motion of the foot rigid body. This rotational component becomes particularly significant during heel-strike or toe-strike events, where small ankle rotations may generate large localized touchdown velocities at the foot edges despite relatively small translational motion of the foot center.

Figure 1. Foot corner contact modeling diagram.

The proposed foot corner contact modeling offers several advantages for learning quiet locomotion. By incorporating a penalty in the subsequent reward design on the foot contact point with the highest impact velocity, a dynamic foot contact penalty is achieved. By constraining extreme local corner velocities rather than global foot motion, the method effectively suppresses high-impact touchdown behaviors and encourages smoother, more controlled foot descent. In addition, the proposed method addresses the transition from point-based foot–ground contact in quadruped robots to surface contact in humanoid robots. By further regulating the velocities of the four foot corners, it enables a coordinated planar touchdown of the foot.

2.2. *Training framework*

To address the challenges arising from partial observability of the environment, an end-to-end reinforcement learning framework based on Proximal Policy Optimization (PPO) is adopted to train the humanoid robot control policy. The overall architecture consists of three core components: an Estimator Encoder, an Actor Policy, and a Critic Network. The training framework is shown in Figure 2.

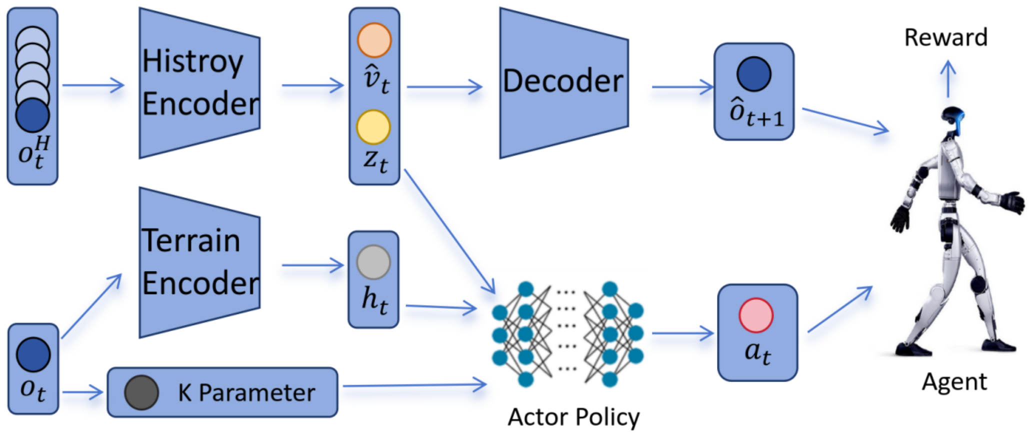

Figure 2. Training framework.

(1) Estimator Encoder: To enhance the observability of the system's dynamic states, an Estimator Encoder is introduced to encode a sequence of historical observations. Given a fixed-length observation window, the Estimator Encoder employs a multi-layer perceptron (MLP) to extract informative features from the observation history $o_t^H = [o_{t-T+1}, \ldots, o_t]$ and outputs two key latent variables: a velocity estimate $\hat{v}_t$ and a dynamic latent vector z_t. The estimator encoder consists of two hidden layers, each with 256 neurons, and employs the Exponential Linear Unit (ELU) activation function in both layers. The observation o_t is derived from the robot's proprioceptive sensors and is defined as:

$$o_t = [\omega_t, g_t, q_t, \dot{q}_t, a_{t-1}, c_{cmd}] \tag{3}$$

where $\omega_t$denotes the angular velocity, g_t represents the gravity vector, q_t and $\dot{q}_t$ correspond to the joint angles and velocities, a_{t-1} is the action taken in the previous timestep, and c_{cmd} represents the command input. Velocity Estimation Loss is defined as follows:

$$L_{est} = \mathbb{E}_t \left[\|\hat{v}_t - v_t\|_2^2 \right] \tag{4}$$

(2) Actor Policy: The Actor Policy is implemented using a MLP. The actor policy network consists of three hidden layers: the first layer has 512 neurons, while the remaining two layers each have 256 neurons. All three layers employ the ELU activation function. Its input is formed by concatenating three components: the current observation o_t, the latent representations produced by the Estimator Encoder (the velocity estimate $\hat{v}_t$ and the dynamic latent vector z_t), and optional terrain encoding features, when height-map information $\hat{h}_t$ is available from the environment. The output action $a_t \in R^{12}$ represents the desired joint position offsets relative to the constant standing joint positions q^{stand}. Consequently, the final desired joint positions are given by:

$$q_t^{des} = q^{stand} + a_t * a_{scale} \tag{5}$$

where a_{scale} is a scaling factor. The loss function of the policy network is defined as:

$$L_{policy} = -\mathbb{E}_t \left[min(r_t(\theta)\hat{A}_t, clip(r_t(\theta), 1-\varepsilon, 1+\varepsilon)\hat{A}_t) \right] \tag{6}$$

where $r_t(\theta)$ represents the importance sampling ratio, ε is a parameter used to suppress overly aggressive updates of the policy network and $\hat{A}_t$ represents the Generalized Advantage Estimation, they are defined as follows:

$$r_t(\theta) = \frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{old}}(a_t|s_t)} \tag{7}$$

$$\hat{A}_t = \sum_{l=0}^{\infty} (\gamma\lambda)^l \delta_{t+1}, \delta_t = r_t + \gamma V_\psi(s_{t+1}) - V_\psi(s_t) \tag{8}$$

where $\pi_\theta(a_t|s_t)$ denotes the policy output by the policy network, and V_ψ represents the value function, which estimates the expected cumulative future return.

(3) Critic Network: The Critic Network is implemented as an independent MLP to approximate the value function $v_\psi(s)$, which estimates the expected return of the current state and supports advantage estimation and value regression in PPO. The critic Network consists of three hidden layers: the first layer has 512 neurons, while the remaining two layers each have 256 neurons. All three layers employ the ELU activation function. During training, the Critic is allowed to leverage richer state information, such as privileged information available in simulation, to improve the accuracy of value learning and thereby stabilize policy optimization. Value loss, and entropy regularization term are defined as:

$$L_{value} = \mathbb{E}_t \left[(V_\psi(s_t) - R_t)^2 \right] \tag{9}$$

$$L_{entropy} = -\mathbb{E}_t \left[\frac{1}{2} \sum_i (1 + log(2\pi\sigma_i^2)) \right] \tag{10}$$

The overall loss function is defined as follows:

$$L_{PPO} = L_{policy} + c_v L_{value} + c_e L_{entropy} + c_{est} L_{est} \tag{11}$$

where c_v, c_e and c_{est} are hyperparameters.

2.3. Reward

As discussed earlier, foot–ground impact is the dominant source of locomotion noise in humanoid robots. Accordingly, the reward function in this study is designed to reduce the touchdown velocity, thereby mitigating locomotion noise. Specifically, the reward imposes constraints on the landing velocity, with stronger penalties applied as the foot approaches the ground. By constraining excessive downward motion when the foot approaches the ground, the proposed pre-contact velocity modulation reduces impact generation at its source. This strategy is consistent with human walking behavior, where active deceleration is applied before foot placement to achieve quiet and compliant contact.

In addition to contact-related objectives, conventional reward terms are incorporated to ensure stable locomotion, accurate command tracking, and natural joint motion. By balancing these objectives, the proposed method enables the policy to achieve quiet locomotion without sacrificing fundamental locomotion performance. The main reward settings are shown in Table 1.

Table 1. Reward setting.

Category	Reward	Definition	Scale [-]
Task Reward	Linear velocity tracking	$e^{-4(v_x^{cmd}-v_x)^2\}}$	1.5
	Linear velocity tracking	$e^{-4(v_y^{cmd}-v_y)^2\}}$	1.0
	Angular velocity tracking	$e^{-4(\omega_z^{cmd}-\omega_z)^2\}}$	2.5
	Base height	$e^{-4(h^{cmd}-h)^2\}}$	2.5
Penalty Reward	Joint torques	$\|\|\tau\|\|^2$	-2.5×10^{-6}
	Base linear velocity Z	$\|\|v_z\|\|^2$	-0.5
	Base orientation	$\|\|\theta_{xy}\|\|^2$	-1.5
	Base angular velocity	$\|\|\omega_{xy}\|\|^2$	-0.025
	Self-collisions	$n_{\text{collision}}$	-1.0
	Foot slip	$\|\|v_{f,xy}\|\|$	-0.25
	Foot air time	$\sum_i(t_{f,air}-0.5)^2$	0.05
Noise Penalty	Drop Foot velocity	$\|\|v_{max}\|\|^2\cdot w_{weight}\cdot K$	-4

2.4. Virtual compliance parameter

In real-world scenarios, humanoid robots face varying requirements on locomotion noise. Quiet and compliant foot–ground contacts are essential in service or indoor environments, while higher impact can be tolerated when efficiency or responsiveness is prioritized. Therefore, a single locomotion policy with fixed “quietness” is insufficient for diverse tasks. To address this, we introduce an adjustable quiet-locomotion control variable that enables dynamic regulation of touchdown compliance within a unified policy.

Inspired by joint impedance control, we aim to endow the foot touchdown behavior with spring-like characteristics. Specifically, as the foot approaches the ground, the vertical touchdown velocity should be minimized, whereas when the foot is farther from the ground, the descent velocity can remain largely unconstrained, allowing for rapid foot motion. To achieve this behavior, the penalty imposed on the foot velocity is coupled with the foot height above the ground, leading to the following height-dependent velocity penalty formulation:

$$penalty = -\|v_{max}\|^2 * w_{weight} * K \tag{12}$$

where, v_{max} denotes the maximum touchdown velocity among the four foot corner points, $w_{weight} = (1 - h_{foot}/h_{max})^2$ is a height-dependent weighting factor, h_{foot} represents the foot height above the ground, and h_{max} is the maximum allowable foot clearance. The parameter K, referred to as the virtual compliance parameter, acts as an equivalent stiffness coefficient of a virtual spring, and is constrained within the range $K \in [0, 1]$.

The virtual compliance parameter K is introduced as a fifth command input, in addition to the three velocity commands and the base height command, to regulate the level of quiet locomotion. A larger value of K corresponds to quieter robot motion. The parameter K is resampled from the range [0,1] at the beginning of each training episode when the robot is reset. This episodic randomization prevents abrupt increases in the loss function during training and contributes to more stable policy optimization.

3. Results

This section evaluates the proposed quiet locomotion control method through comprehensive simulation and real-world experiments. Baseline controller and controller used for ablation study are compared to analyze foot touchdown dynamics and locomotion noise characteristics. Quantitative metrics on foot descending velocity and acoustic noise are used to assess the effectiveness of pre-contact velocity modulation and foot corner contact modeling. Both simulation and real world results demonstrate that the proposed approach significantly reduces acoustic noise while preserving stable and robust walking performance, and achieves compliant contact.

3.1. Experimental setup

(1) Simulation Training: Reinforcement learning training is conducted using the NVIDIA Isaac Gym physics simulation platform, with the Unitree G1 humanoid robot model as the training subject. During training, the robot performs velocity-tracking locomotion tasks on flat terrain, while multiple reward terms are employed to constrain stability, motion smoothness, and foot–ground contact behaviors. In the simulation environment, the frequency of the policy output is 50 Hz, while the low-level torque output operates at 200 Hz. All training experiments are executed on a single NVIDIA RTX 4060 GPU, with the complete training process taking approximately 12 hours. This setup ensures efficient training while demonstrating the feasibility of the proposed method under moderate computational resources.

(2) Real-World Experiments: After simulation training is completed, the learned policy is directly transferred to the real Unitree G1 robot without further fine-tuning. The robot relies exclusively on proprioceptive sensing, including joint positions, joint velocities, and body orientation. Experiments are conducted on an indoor flat surface. In the real-world environment, the frequency

of the policy output is 50 Hz, while the low-level torque output operates at 200 Hz. To evaluate the acoustic noise of the locomotion process, the robot is additionally equipped with a sound pressure sensor to collect noise generated during walking. The sound pressure sensor was mounted near the robot waist to measure the overall locomotion noise generated during walking. Placing the sensor near the feet would expose it to strong local disturbances caused by vigorous foot swinging motions and direct foot-ground collisions during touchdown. These impact-induced structural vibrations could significantly interfere with the acoustic measurements and introduce large fluctuations unrelated to the overall locomotion noise characteristics. Therefore, the waist position was selected to reduce the influence of localized impact disturbances and obtain more stable measurements of the global airborne and structural noise generated during robot locomotion. The sound pressure sensor internally performs acoustic signal acquisition at a sampling frequency of 16 kHz for broadband noise measurement. For locomotion noise evaluation and synchronization with robot motion data, the processed sound pressure level (SPL) values were recorded at an update rate of 10 Hz. The experimental setup in the real-world experiments is illustrated in Figure 3.

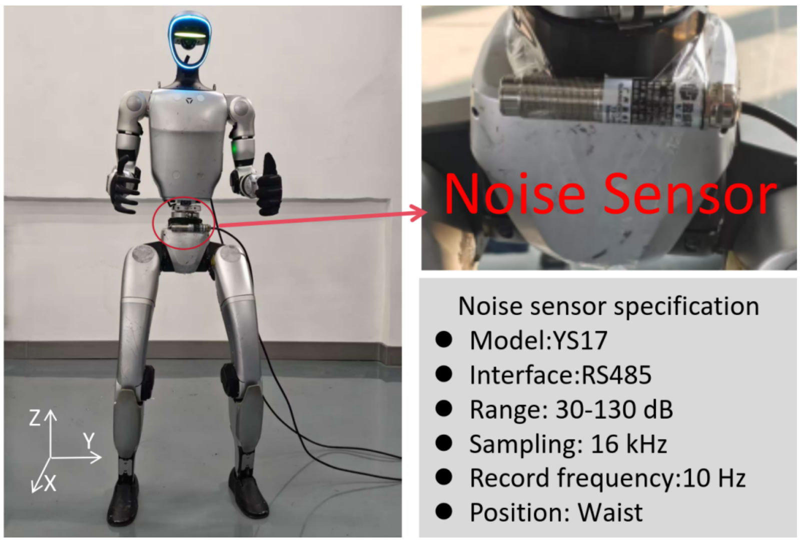

Figure 3. Experimental setup for real-world.

In both simulation and real-world experiments, four control strategies were evaluated. These include: (1) the reinforcement learning–based control policy provided by Homie [29] as the baseline, which can train a Unitree G1 robot to walk and squat under arbitrary continuous changing upper-body poses; (2) an ablation method in which the foot corner contact modeling is removed, the parameter K is fixed at 1, corresponding to the maximum quietness level; and (3–4) two versions of the proposed method with the virtual compliance parameter set to 0 and 1, respectively.

To ensure the consistency and fairness of the locomotion noise evaluation, all real-world experiments were conducted under identical experimental conditions. Specifically, the proposed method and baseline method were tested in the same indoor environment (a corridor) using the same floor surface (smooth marble flooring) and sensor placement (robot waist). The sound pressure sensor position remained fixed throughout all experiments to reduce measurement variability caused by environmental or positional changes.

3.2. Simulation results

In simulation, four control strategies were evaluated within the Isaac Gym environment. All methods are commanded to track the same forward velocity of 0.4 m/s and are evaluated on flat terrain to ensure a fair and consistent comparison. Since noise cannot be directly modeled in the simulation environment, the foot velocity is used to reflect the level of quiet locomotion. At each time step, the largest negative Z-axis velocity among the four foot contact points is defined as the falling velocity. We then compute the temporal average of the falling velocity, as well as the downward velocity at the instant of contact, as summarized in Table 2. To ensure the continuity of the velocity profile, we use the velocity of the center of mass to plot the velocity curves of the four methods, as shown in Figure 4.

Table 2. Statistics of the maximum vertical (Z-axis) descending velocity among the four foot contact points for four methods.

Method	Average velocity (m/s)		Contact velocity (m/s)	
	Left foot	**Right foot**	**Left foot**	**Right foot**
Baseline	−0.292	−0.248	−0.751 ± 0.264	−0.504 ± 0.219
Ablation	−0.407	−0.240	−0.908 ± 0.198	−0.581 ± 0.077
K = 0	−0.208	−0.102	−0.396 ± 0.063	−0.251 ± 0.049
K = 1	−0.202	−0.092	−0.406 ± 0.079	−0.229 ± 0.039

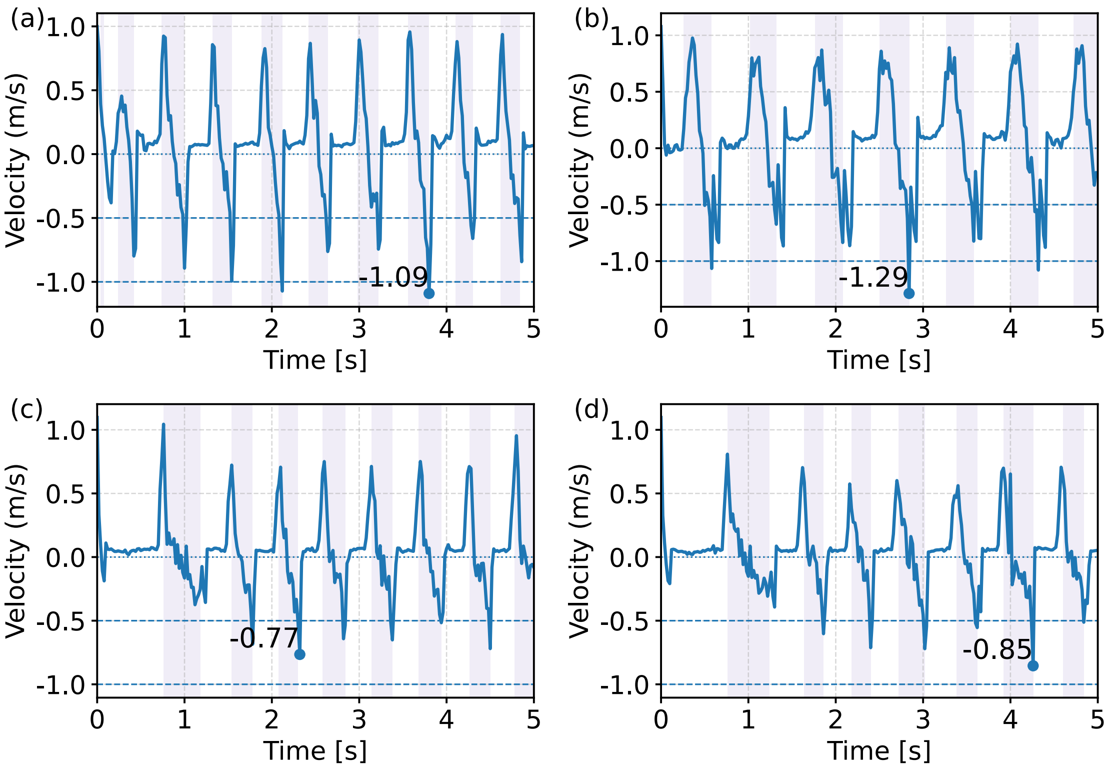

Figure 4. Left-Foot vertical velocity for the four methods during 5 s of steady-state walking. The light purple shaded regions denote the foot swing phase, whereas the white regions indicate periods of ground contact. **(a)** Baseline Homie; **(b)** Ablation without Foot Corner Contact Modeling; **(c)** Proposed Method K = 0; **(d)** Proposed Method K = 1.

As indicated by the results in Table 2, the proposed method significantly reduces both the foot descending velocity and the velocity at ground contact, which effectively contributes to noise reduction during locomotion and enables more compliant contact. In addition, even when the virtual compliance parameter is set to zero, the proposed approach still achieves a noticeable reduction in descending velocity compared to the baseline method and ablation method. This behavior is attributable to the command randomization employed during training, which encourages the robot to learn quieter walking behaviors even in the absence of explicit virtual compliance parameter.

The ablation strategy that solely characterizes foot movement based on the velocity of the foot's center of mass is difficult to capture the local contact process, and can easily trigger a landing pattern of "toe first touching the ground followed by rapid heel tapping," resulting in greater instantaneous ground contact velocity. In contrast, the proposed method maintains both swing-phase and touchdown velocities within a lower and safer range, resulting in smoother and quieter locomotion. When K = 1, both the average descending velocity and the contact velocity are consistently lower than those obtained with K = 0, demonstrating that the policy has successfully learned to modulate different levels of quiet walking behavior.

To evaluate the robot's velocity tracking capability, the velocity tracking curves of the robot are plotted as shown in Figure 5. As shown in the velocity tracking results, the proposed method exhibits a moderate reduction in tracking responsiveness compared with the baseline controller. This behavior is mainly caused by the compliant touchdown strategy introduced for impact suppression, which intentionally reduces touchdown velocity and contact intensity during gait transitions. Consequently, the robot responds more conservatively to commanded motion changes. However, despite the slight degradation in velocity tracking performance, the robot maintains stable locomotion behavior throughout the simulation without observable gait instability or loss of balance. The results indicate that the proposed controller achieves quieter and softer touchdown behavior while preserving overall locomotion stability.

Considering the magnitude of the contact force, we further record the peak plantar contact forces of the robot during simulation. Given that the G1 robot has a body mass of approximately 35 kg, contact forces exceeding 330 N are defined as impact contact forces. For each of the four methods, the average value of these impact contact forces is computed and summarized in the Table 3. In addition, the right-foot contact force profiles of the four methods are presented for comparison, illustrating the temporal evolution of ground reaction forces during locomotion.

As shown in Figure 6, selecting contact forces greater than 330 N effectively characterizes the impact events occurring at touchdown. Furthermore, the results presented in Table 3 indicate that the proposed method significantly reduces the magnitude of contact forces at the moment of landing. This reduction demonstrates improved compliant control performance, which consequently contributes to mitigating locomotion-induced noise.

To further investigate the impact characteristics at touchdown, the ankle joint torque and angular velocity are analyzed. The G1 humanoid robot is equipped with two ankle joints per foot, namely a pitch joint and a roll joint. Since the roll joint primarily contributes to lateral stability during stance, while vertical impact mitigation is mainly governed by the pitch joint, the present analysis focuses on the pitch joint dynamics.

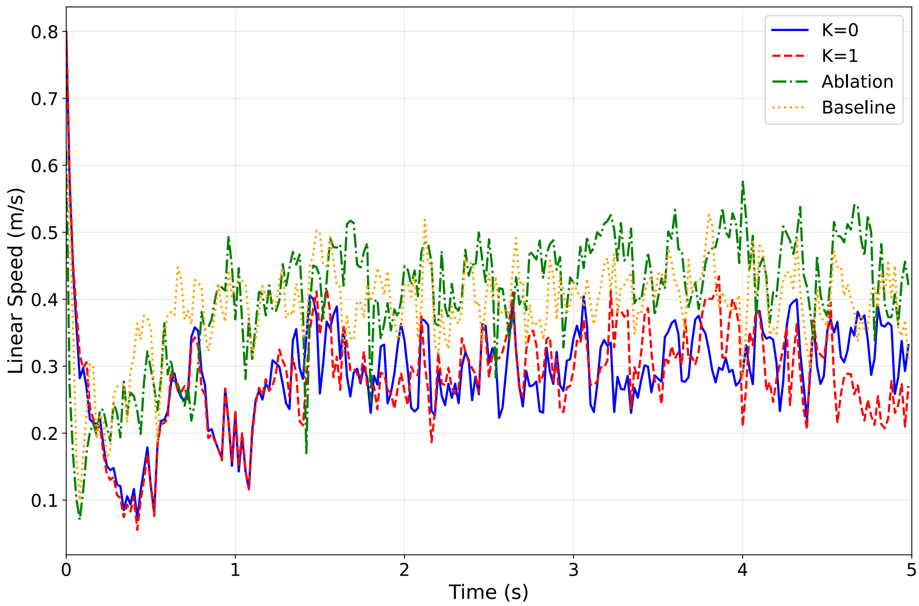

Figure 5. Comparison of velocity tracking curves of four methods.

Table 3. Average and peak impact contact forces for four methods.

Method	Average force (N)		Peak force (N)	
	Left foot	Right foot	Left foot	Right foot
Baseline	485.52	510.94	1401.83	2160.93
Ablation	425.09	687.55	1997.78	2778.95
K = 0	424.53	407.04	826.64	801.94
K = 1	400.56	397.44	614.77	611.64

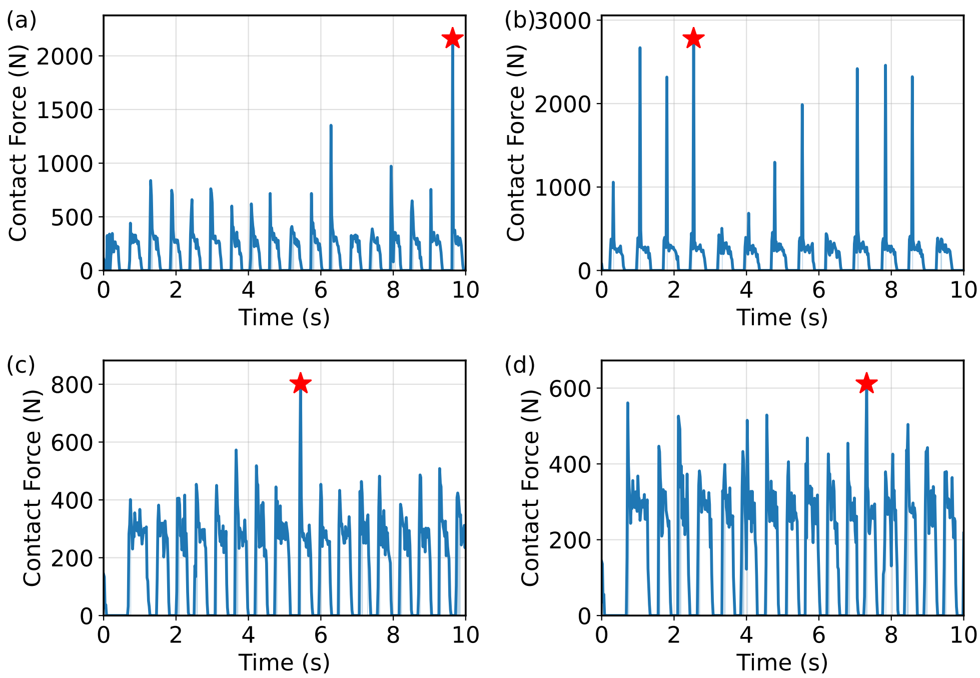

Figure 6. Contact force profiles of the right foot across four methods during stable walking. **(a)** Baseline Homie; **(b)** Ablation without Foot Corner Contact Modeling; **(c)** Proposed Method K = 0; **(d)** Proposed Method K = 1.

For each of the four methods, the peak values of ankle pitch torque and angular velocity for both the left and right feet are recorded. The corresponding quantitative results are summarized in Table 4. In addition, comparative plots of the pitch joint torque and angular velocity profiles are presented in Figure 7 and Figure 8, respectively, to illustrate the dynamic differences among the four methods.

As shown in Table 4, the ablation method exhibits relatively smaller negative torque peaks and larger positive torque peaks. This indicates the absence of an active toe-lifting control prior to touchdown, resulting in increased impact at contact. This observation is consistent with the contact pattern of the ablation method, where the toe makes contact first, followed by a relatively large heel impact. In contrast, the remaining three methods show comparable torque profiles, which are generally consistent with the previously observed trends in touchdown velocity and ground contact force.

Table 4. Peak foot ankle pitch velocity and torque for four methods.

Method	Peak Velocity + (rad/s)		Peak Velocity − (rad/s)		Peak Torque + (N·m)		Peak Torque − (N·m)	
	Left	Right	Left	Right	Left	Right	Left	Right
Baseline	14.5545	14.7811	−10.0115	−4.5187	27.8880	28.1824	−22.7928	−19.3467
Ablation	5.7670	6.6185	−13.0785	−12.0611	37.6593	34.6323	−8.6262	−3.7506
K = 0	8.5427	7.5077	−3.3778	−3.8678	27.2310	25.9786	−21.8193	−16.6931
K = 1	7.9030	6.8043	−3.3778	−3.8678	28.8923	26.5631	−24.4130	−14.0466

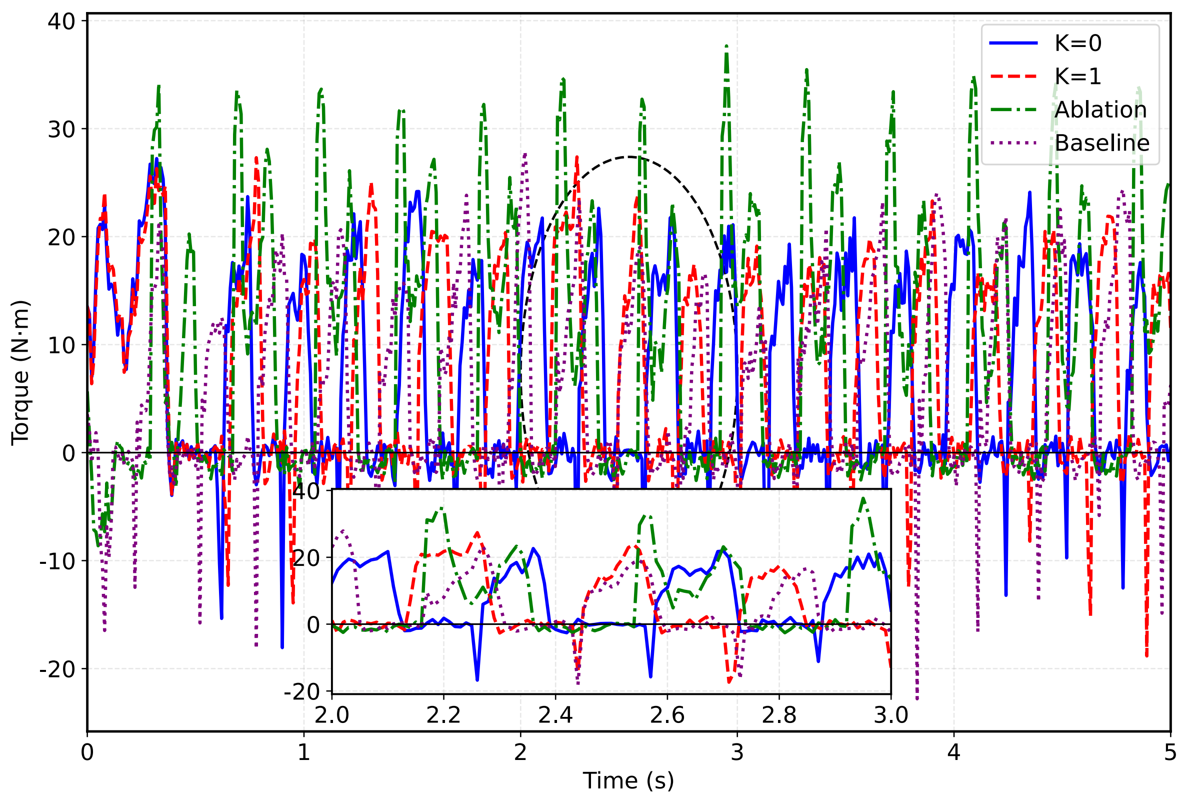

Figure 7. Left foot ankle pitch joint torque for four methods (0–5 s).

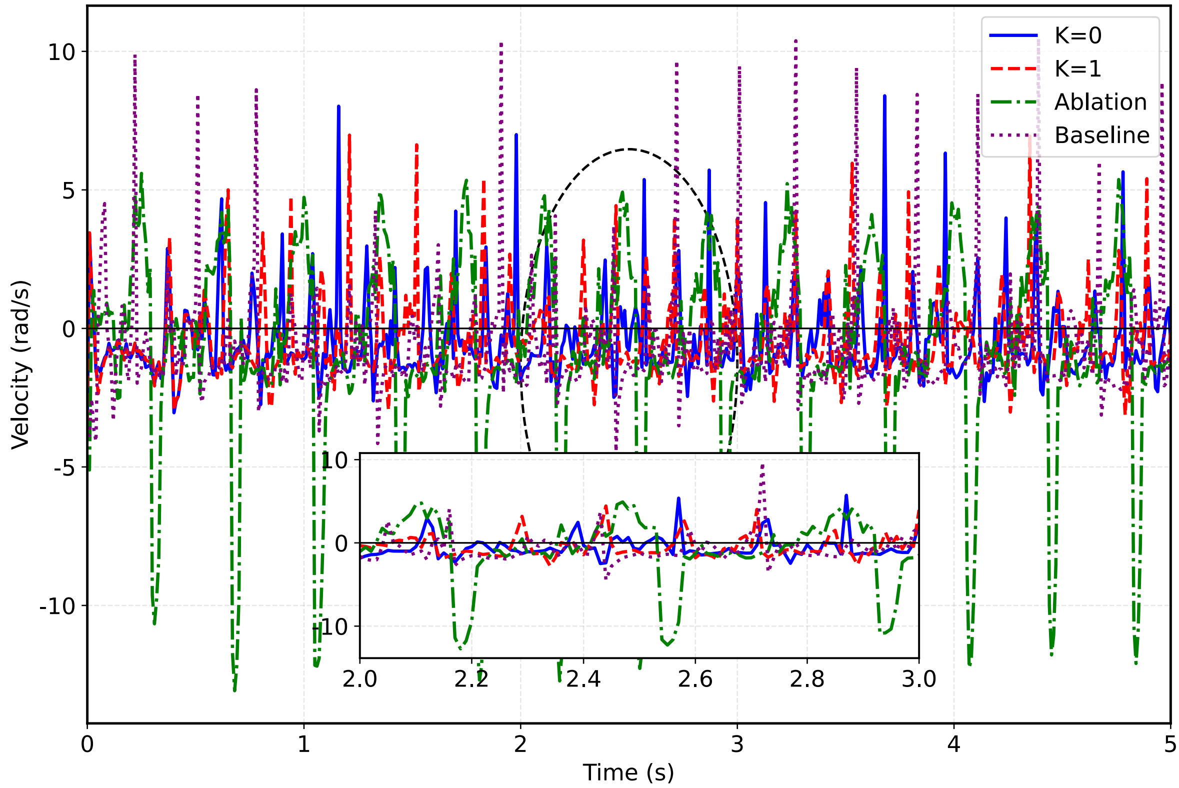

Figure 8. Left foot ankle pitch joint velocity across four methods (0–5 s).

Compared to joint torque, joint angular velocity reflects the combined effect of actively applied torque and external ground impact, and thus provides a more direct indicator of impact severity. As shown in the table, the proposed method yields smaller negative angular velocity peaks, corresponding to a smoother and more compliant ground contact. In comparison, both the baseline and the ablation methods exhibit larger negative angular velocity peaks, indicating rapid post-impact posture adjustments caused by stronger ground collisions. These results demonstrate that the proposed method achieves a more compliant and impact-mitigated landing behavior.

The peak torques of all leg joints under the four methods are summarized in Table 5. As shown in the Table 5, the proposed method effectively reduces peak torques across almost all joints, especially at the knee joints. For the right leg, compared with the baseline method, the peak torque is reduced by 19.88 N·m. These results demonstrate that the proposed approach enables compliant landing, which can significantly mitigate hardware stress and potential damage to the robot.

To investigate the effectiveness of different control strategies under varying commanded speeds, a comparative study is conducted, while all other training settings remain unchanged. The resulting policies are evaluated alongside the baseline over commanded forward velocities ranging from 0.2 m/s to 0.6 m/s. The corresponding quantitative results are summarized in Table 6.

As evidenced by the quantitative results in Table 6, the proposed method exhibits stable and controlled touchdown behavior under all commanded speeds. In contrast, the strategy that represents foot motion solely using the foot center-of-mass velocity shows increasing sensitivity to speed variations. Without explicitly modeling localized contact dynamics, it tends to produce a toe-first contact pattern followed by a rapid heel impact, particularly at higher commanded speeds, resulting in larger instantaneous contact velocities. Consequently, across the velocity range from 0.2 m/s to 0.6 m/s, the touchdown velocities of the ablated model remain consistently higher than those achieved by the proposed method and the Homie baseline, indicating inferior robustness to speed escalation.

Table 5. Peak torques of all leg joints for four methods.

Joint	Baseline (N·m)	Ablation (N·m)	K = 0 (N·m)	K = 1 (N·m)
L_Hip_Pitch	33.64	39.60	37.38	36.58
L_Hip_Roll	30.90	36.36	35.99	33.82
L_Hip_Yaw	26.67	26.97	23.95	24.22
L_Knee	57.85	59.03	45.84	50.41
L_Ankle_Pitch	27.89	37.66	27.23	28.89
L_Ankle_Roll	24.06	26.20	20.48	25.86
R_Hip_Pitch	30.21	40.42	38.69	40.73
R_Hip_Roll	24.63	31.70	42.81	43.00
R_Hip_Yaw	30.81	31.11	29.05	28.85
R_Knee	64.52	61.21	44.64	44.60
R_Ankle_Pitch	28.18	34.63	25.98	26.56
R_Ankle_Roll	27.78	23.96	18.39	18.95

Table 6. Comparison of foot descending velocities under different methods and commanded speeds.

Method	Command (m/s)	Average velocity (m/s)		Contact velocity (m/s)	
		Left	Right	Left	Right
Baseline	0.2	−0.113	−0.060	−0.290 ± 0.098	−0.181 ± 0.062
	0.3	−0.220	−0.134	−0.553 ± 0.201	−0.297 ± 0.094
	0.4	−0.292	−0.248	−0.751 ± 0.264	−0.504 ± 0.219
	0.5	−0.361	−0.259	−0.729 ± 0.236	−0.467 ± 0.153
	0.6	−0.422	−0.289	−0.798 ± 0.242	−0.529 ± 0.251
Ablation	0.2	−0.263	−0.194	−0.627 ± 0.195	−0.384 ± 0.137
	0.3	−0.338	−0.223	−0.775 ± 0.181	−0.557 ± 0.133
	0.4	−0.407	−0.240	−0.908 ± 0.198	−0.581 ± 0.077
	0.5	−0.450	−0.280	−0.967 ± 0.161	−0.639 ± 0.081
	0.6	−0.489	−0.339	−1.026 ± 0.166	−0.724 ± 0.111
K = 0	0.2	−0.115	−0.063	−0.286 ± 0.069	−0.156 ± 0.045
	0.3	−0.167	−0.086	−0.353 ± 0.080	−0.204 ± 0.038
	0.4	−0.208	−0.102	−0.396 ± 0.063	−0.251 ± 0.049
	0.5	−0.240	−0.120	−0.440 ± 0.076	−0.311 ± 0.055
	0.6	−0.260	−0.145	−0.479 ± 0.094	−0.379 ± 0.061
K = 1	0.2	−0.083	−0.061	−0.214 ± 0.047	−0.147 ± 0.034
	0.3	−0.133	−0.075	−0.293 ± 0.067	−0.170 ± 0.043
	0.4	−0.202	−0.092	−0.406 ± 0.079	−0.229 ± 0.039
	0.5	−0.226	−0.108	−0.426 ± 0.110	−0.271 ± 0.048
	0.6	−0.259	−0.132	−0.506 ± 0.128	−0.362 ± 0.063

Although all methods exhibit increasing touchdown velocities as the commanded speed increases, the proposed method consistently maintains the lowest descending and contact velocities throughout the entire speed range, effectively mitigating impact forces and locomotion noise. At low commanded speeds, the differences among methods become less pronounced due to the reduced velocity tracking demands. Furthermore, as shown in Figure 5, the actual tracking capability of the proposed method at a commanded speed of 0.4 m/s is 0.3 m/s. Compared with the baseline and ablation methods at a commanded speed of 0.3 m/s, the landing speed of both feet is still significantly reduced, demonstrating the effectiveness of our method under the same speed tracking capability.

3.3. Real world results

To validate the effectiveness of the proposed method, real-world experiments were conducted on a humanoid robot and compared with the baseline method and ablation method. All real-world experiments were conducted in a relatively open indoor laboratory environment with smooth marble flooring and no dedicated acoustic treatment. To ensure fair comparison, the proposed method and baseline method were evaluated under identical environmental conditions, including the same room layout, floor surface, robot hardware configuration, walking commands, and sensor placement. For all methods, the commanded forward velocity was set to 0.4 m/s, and each trial required the robot to walk continuously for at least 10 seconds. Five independent trials were conducted for each method.

The ambient noise level measured when the robot remained stationary was approximately 62 dB. The definition of decibels is as follows: $L = 10\lg(I/I_0)$, where I represents the sound energy. For each trial, both the peak noise level and the average noise level during locomotion were recorded. After converting the measured sound pressure levels into acoustic energy and removing the background noise component, the resulting change in the measured locomotion noise was relatively small due to the large energy difference between the locomotion noise and ambient noise levels. Therefore, all experiments directly used the measured noise as the motion-generated noise. The reported results correspond to the mean values of the peak noise and the mean noise level over the five trials, as summarized in Table 7.

Table 7. Comparison of average and peak noise levels for four methods.

Method	Average noise (dB)	Peak noise (dB)
Baseline	87.24 ± 1.10	97.98 ± 1.92
Ablation	87.41 ± 1.92	110.52 ± 0.88
K = 0	82.99 ± 1.03	95.00 ± 2.40
K = 1	82.36 ± 0.76	93.10 ± 2.30

The results presented in Table 7 and Figure 9 demonstrate that the proposed method can significantly reduce walking-induced noise compared with the baseline approach. Under the K = 1 command condition, the average noise level is reduced by 4.88 dB, while the peak noise level is reduced by 4.88 dB. In comparison with the ablation method, the proposed approach achieves a reduction of 5.05 dB in average noise and 17.42 dB in peak noise. It is worth noting that the noise sensor is mounted on the robot's waist, relatively close to the feet. Since acoustic intensity attenuates rapidly with increasing distance, the perceived noise reduction at a normal human–robot interaction distance is expected to be more pronounced than the measured 5 dB reduction.

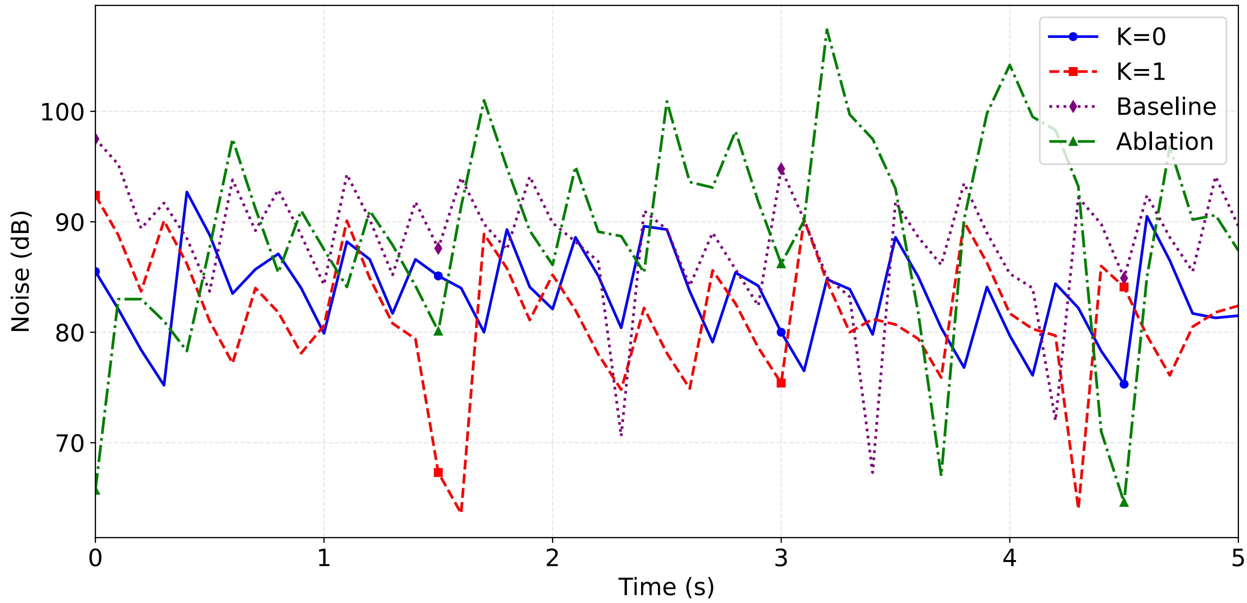

Figure 9. Comparison of noise across four methods during five seconds of steady-state locomotion.

When K = 0, a reduction in walking noise is also observed compared with the baseline method; however, the noise suppression effect achieved with K = 1 is consistently superior. This indicates that the robot has learned different degrees of quiet walking behavior, and that the task command variable K enables continuous modulation of walking quietness, allowing the level of noise reduction to be adjusted according to task requirements.

Considering the impact of the method on speed tracking, we ran the baseline method with a speed command of 0.3 m/s to achieve the same actual speed as the proposed method. Under otherwise identical conditions, five experiments were conducted, and the measured average noise was 85.62 dB. Our method still achieved a reduction of 3.26 dB. The experimental results demonstrate the effectiveness of our method under the same actual speed conditions.

To further analyze the acoustic characteristics of the proposed method, frequency-domain analysis was conducted based on the recorded locomotion noise signals. Figure 10 shows the power spectral density (PSD) comparison between the baseline method and the proposed method, while Figure 11 presents the absolute energy distribution across different frequency bands.

From the PSD results, it can be observed that the proposed method reduces acoustic energy over a broad frequency range. In particular, a pronounced reduction is found in the low-frequency region, which is commonly associated with foot–ground impact events during locomotion. These low-frequency components correspond to the perceptible “thumping” sounds generated at touchdown. The band-energy analysis further confirms that a substantial portion of the noise reduction is concentrated in the low-frequency bands. Since low-frequency components are dominant in humanoid locomotion noise, their suppression can significantly reduce the perceived walking noise. In addition, reductions are observed across other frequency bands. This is because foot–ground impacts can be transmitted through the mechanical structure, exciting structural vibrations and inducing additional higher-frequency noise components. By reducing the impact intensity at touchdown, the vibration-induced noise associated with these transmitted impacts is also attenuated.

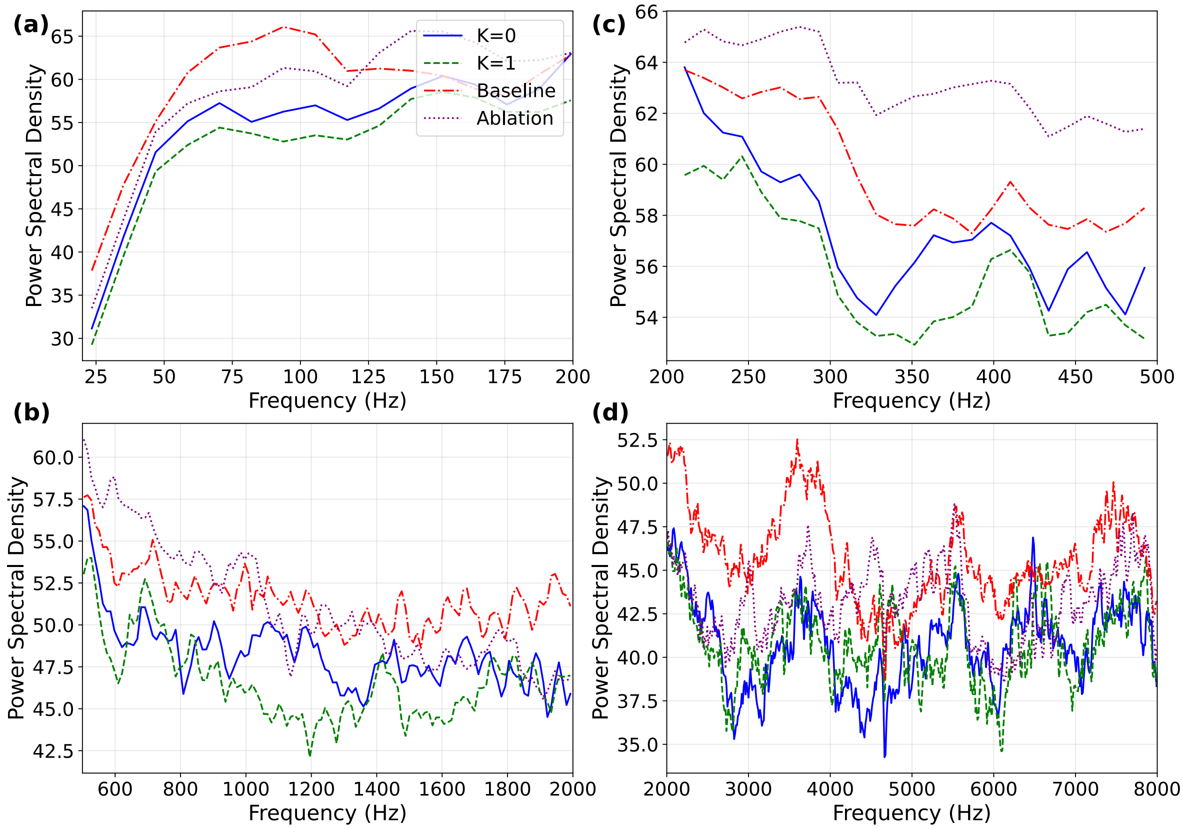

Figure 10. Comparison of power spectral density for four methods at different frequencies. **(a)** 20–200 Hz; **(b)** 200–500 Hz; **(c)** 500–2000 Hz; **(d)** 2000–8000 Hz.

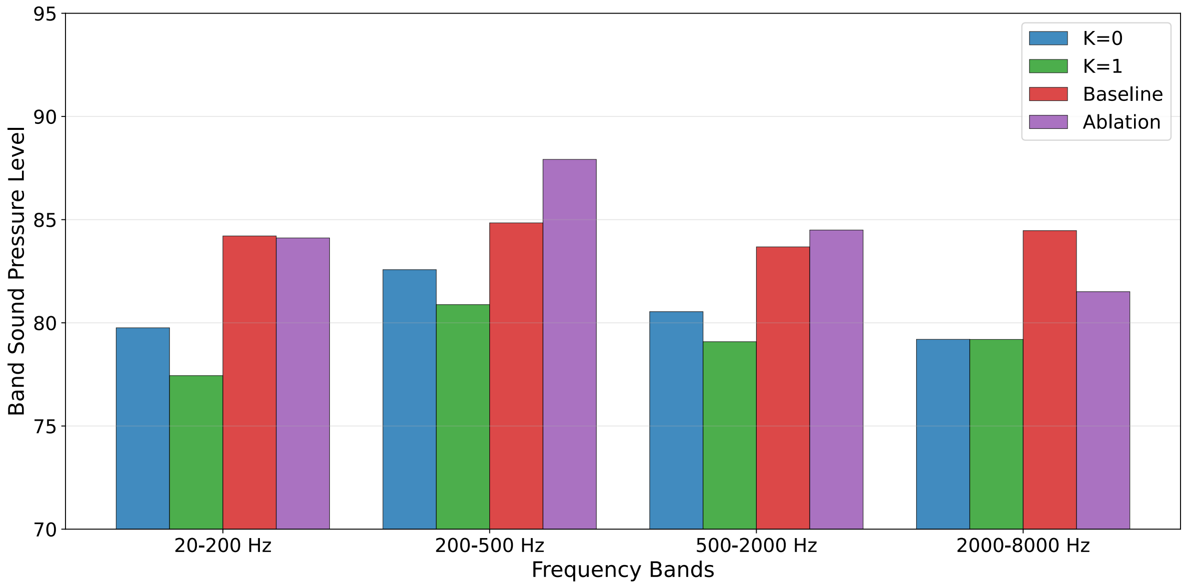

Figure 11. Comparison of acoustic energy for four methods.

Considering the ankle pitch joint velocity and torque, the peak values of the joint velocity and torque over the five trials are obtained and summarized in Table 8. The velocity and torque curves of the left ankle joint during five seconds of stable walking are illustrated in Figure 12 and Figure 13, respectively.

Table 8. Peak ankle pitch velocity and torque for four methods.

Method	Peak Velocity + (rad/s)		Peak Velocity − (rad/s)		Peak Torque + (N·m)		Peak Torque − (N·m)	
	Left	Right	Left	Right	Left	Right	Left	Right
Baseline	7.06 ± 2.03	7.57 ± 1.08	−1.35 ± 0.06	−1.65 ± 0.14	31.98 ± 6.49	28.32 ± 2.85	−18.57 ± 4.21	−19.49 ± 2.05
Ablation	5.05 ± 0.76	4.91 ± 0.27	−5.55 ± 0.87	−4.95 ± 0.13	44.12 ± 7.37	41.32 ± 4.79	−4.99 ± 3.87	−5.31 ± 5.00
K = 0	4.60 ± 0.78	4.82 ± 3.49	−1.29 ± 0.28	−1.86 ± 2.23	33.88 ± 8.35	30.24 ± 2.89	−12.37 ± 5.38	−8.71 ± 4.16
K = 1	4.60 ± 0.30	3.24 ± 0.73	−1.46 ± 0.17	−1.25 ± 0.16	29.97 ± 0.55	28.30 ± 2.17	−12.03 ± 2.78	−7.09 ± 1.88

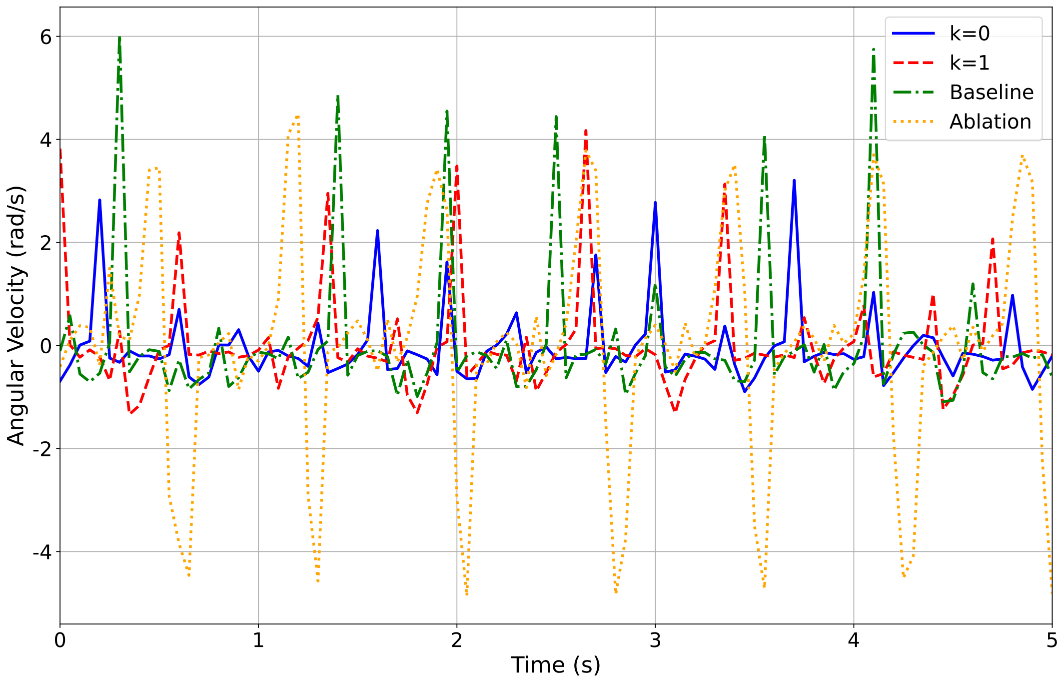

Figure 12. Comparison of left ankle joint velocity during stable walking for four methods.

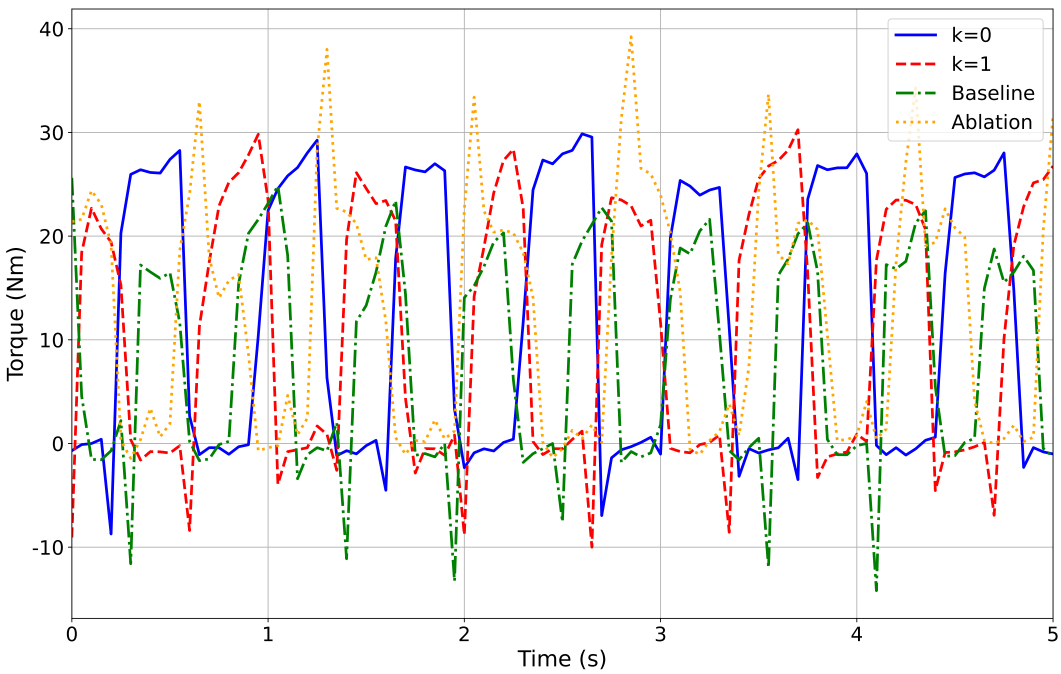

Figure 13. Comparison of left ankle joint torque during stable walking for four methods.

According to the experimental results, the ablation method exhibits the smallest peak value of negative torque. This is because the absence of active toe-lifting regulation causes the robot to land directly on the toe, followed by a heavy heel impact, resulting in the largest positive torque peak and the least compliant landing behavior. This indicates that a smaller impact is generated, which is consistent with the simulation results.

Considering compliant control of the robot, the peak torques of all leg joints are evaluated. The average peak torques over five experiments are summarized in Table 9. As shown in the table, the proposed method significantly reduces the peak torques of almost all leg joints, especially at the knee joints. These results demonstrate the effectiveness of the proposed method in achieving compliant joint control, thereby significantly reducing potential joint damage in the robot. The experimental data demonstrate that the proposed method enables the robot to learn a more compliant landing strategy, significantly reducing the noise generated during walking.

Table 9. Peak torques of all leg joints for different methods.

Joint	Baseline (N·m)	Ablation (N·m)	K = 0 (N·m)	K = 1 (N·m)
L_Hip_Pitch	34.47 ± 4.01	39.12 ± 10.99	35.28 ± 4.24	28.47 ± 2.13
L_Hip_Roll	29.86 ± 2.53	35.66 ± 11.71	35.19 ± 6.63	25.59 ± 7.82
L_Hip_Yaw	24.96 ± 0.71	23.26 ± 3.02	26.52 ± 2.28	25.59 ± 4.90
L_Knee	52.54 ± 3.59	52.13 ± 5.03	46.42 ± 12.64	43.15 ± 7.91
L_Ankle_Pitch	31.98 ± 6.49	44.12 ± 7.37	33.88 ± 8.35	29.97 ± 0.55
L_Ankle_Roll	6.56 ± 1.84	6.69 ± 2.03	7.53 ± 3.31	6.29 ± 1.21
R_Hip_Pitch	36.70 ± 5.35	38.14 ± 5.03	42.76 ± 6.62	33.97 ± 2.77
R_Hip_Roll	35.19 ± 3.04	37.61 ± 2.38	43.40 ± 11.79	37.77 ± 3.69
R_Hip_Yaw	25.12 ± 0.66	24.66 ± 1.40	32.50 ± 5.75	29.20 ± 3.40
R_Knee	50.34 ± 6.08	50.29 ± 4.64	44.96 ± 7.33	41.08 ± 6.55
R_Ankle_Pitch	28.32 ± 2.89	41.32 ± 4.79	30.24 ± 2.89	28.30 ± 2.17
R_Ankle_Roll	4.49 ± 0.68	6.02 ± 2.27	8.94 ± 4.92	7.98 ± 2.99

The trends observed in the real-world experimental results are qualitatively consistent with those observed in simulation. In simulation, the proposed method reduces the foot contact velocity, while in real-world experiments, it reduces the noise generated by the robot. In addition, the peak knee joint torque is significantly reduced in both environments, indicating that the proposed corner-based contact regulation strategy captures physically meaningful touchdown dynamics. This enables a significant reduction in contact impact and achieves compliant landing, rather than merely emulating specific behaviors in simulation.

4. Conclusion

This paper proposes a reinforcement learning–based foot–ground contact modulation method for quiet humanoid locomotion. By incorporating a corner-level foot modeling strategy, the method captures localized contact behaviors and enables effective regulation of pre-contact foot velocity through a dedicated penalty design, thereby suppressing excessive impact during touchdown. Furthermore, a virtual compliance parameter is introduced to continuously adjust the level of walking quietness within a unified control framework. Compared with baseline and ablation methods, the proposed approach significantly reduces foot touchdown velocity and peak contact forces in simulation. It is further validated on a real humanoid robot, reducing the average locomotion noise by 4.88 dB, demonstrating its effectiveness for quiet humanoid locomotion and compliant contact in humanoid robots. Future work will focus on extending the proposed method to more complex and dynamic environments, improving robustness under external disturbances, and exploring its integration with perception-driven adaptive locomotion strategies.

Data availability statement

The datasets generated in this study are available on GitHub at https://github.com/jkfly666/quiet_locomotion_data.

Declaration of generative AI and AI-assisted technologies

During the preparation of this manuscript, the authors used generative AI tools (deepseek-v3.0) only to improve language and readability. The authors take full responsibility for the content of the manuscript.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (62303137, Xiangyu Shao; 62303135, Yi Zeng), in part by the PhD Candidate Innovation Fund of State Key Laboratory of Robotics and Systems (SKLRS-2025-ZM-10, Shaojie Zhang), and in part by the Natural Science Foundation of Sichuan Province (2025ZNSFSC0461, Xiangyu Shao). We would also like to thank the Guangdong Provincial Key Lab of Robotics and Intelligent System for insightful discussions about the idea and analysis.

Authors' contribution

Conceptualization, Xiangyu Shao and Kai Jian; methodology, Kai Jian and Xiangyu Shao; data curation, Kai Jian; writing—original draft preparation, Kai Jian; writing—review and editing, Dong Zhu, Shaojie Zhang and Xiangyu Shao; funding acquisition, Shaojie Zhang and Xiangyu Shao. All authors have read and agreed to the published version of the manuscript.

Conflicts of interest

Dong Zhu is employee of Chengdu CHG Robots & Intelligent Equipment Institute Co., Ltd. The author declares that this affiliation does not influence the design, analysis, interpretation, or presentation of the results reported in this manuscript. The other authors declare that they do not have any conflicts of interest.

References

[1] Margolis GB, Yang G, Paigwar K, Chen T, Agrawal P. Rapid locomotion via reinforcement learning. *Int. J. Rob. Res.* 2024, 43(4):572–587.

[2] Lee J, Hwangbo J, Wellhausen L, Koltun V, Hutter M. Learning quadrupedal locomotion over challenging terrain. *Sci. Rob.* 2020, 5(47):eabc5986.

[3] Rudin N, Hoeller D, Reist P, Hutter M. Learning to walk in minutes using massively parallel deep reinforcement learning. In *Proceedings of the 5th Conference on Robot Learning, PMLR*, London, UK, November 8–11, 2022, pp. 91–100.

[4] Jenelten F, He J, Farshidian F, Hutter M. DTC: deep tracking control. *Sci. Rob.* 2024, 9(86):eadh5401.

[5] Wang Q, Wang Y, Wang H. InstaDrive: street view generation based on the unified instance segmentation input of vehicles and map elements. *Robot Learn.* 2026, 3(1):0004.

[6] Ulloa CC, del Cerro J, Barrientos A. Mixed-reality for quadruped-robotic guidance in SAR tasks. *J. Comput. Des. Eng.* 2023, 10(4):1479–1489.

[7] Go E, Lee JH, Kim SY, Lee JS, Kim HH, *et al.* Stakeholder perspectives on safety issues in collaborative mobile robots: a case study of quadruped robot applications in a smart factory. *Appl. Sci.* 2024, 14(22):10232.

[8] Yu P, Dong S, Sheng S, Feng L, Kwiatkowska M. Trust-aware motion planning for human-robot collaboration under distribution temporal logic specifications. In *Proceedings of 2024 IEEE International Conference on Robotics and Automation (ICRA)*, Yokohama, Japan, May 13–17, 2024, pp. 12949–12955.

[9] Wu Z, Zheng K, Ding Z, Gao H. A survey on legged robots: advances, technologies and applications. *Eng. Appl. Artif. Intell.* 2024, 138:109418.

[10] Tong Y, Liu H, Zhang Z. Advancements in humanoid robots: a comprehensive review and future prospects. *IEEE CAA J. Autom. Sinica* 2024, 11(2):301–328.

[11] Hoeller D, Rudin N, Sako D, Hutter M. ANYmal parkour: learning agile navigation for quadrupedal robots. *Sci. Rob.* 2024, 9(88):eadi7566.

[12] Cheng X, Shi K, Agarwal A, Pathak D. Extreme parkour with legged robots. In *Proceedings of 2024 IEEE International Conference on Robotics and Automation (ICRA)*, Yokohama, Japan, May 13–17, 2024, pp. 11443–11450.

[13] Allen A, Drummond T, Kulić D. Sound judgment: properties of consequential sounds affecting human-perception of robots. In *Proceedings of 2025 20th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, Melbourne, Australia, March 4–6, 2025, pp. 94–102.

[14] Tennent H, Moore D, Jung M, Ju W. Good vibrations: how consequential sounds affect perception of robotic arms. In *Proceedings of 2017 26th IEEE international symposium on robot and human interactive communication (RO-MAN)*, Lisbon, Portugal, August 28–September 1, 2017, pp. 928–935.

[15] Trovato G, Paredes R, Balvin J, Cuellar F, Thomsen NB, *et al.* The sound or silence: investigating the influence of robot noise on proxemics. In *Proceedings of 2018 27th IEEE international symposium on robot and human interactive communication (RO-MAN)*, Nanjing, China, August 27–31, 2018, pp. 713–718.

[16] Watanabe R, Miki T, Shi F, Kadokawa Y, Bjelonic F, *et al.* Learning quiet walking for a small home robot. In *Proceedings of 2025 IEEE International Conference on Robotics and Automation (ICRA)*, Atlanta, USA, May 19–23, 2025, pp. 15285–15291.

[17] Cao Z, Nie B, Zhang Y, Gao Y. Minimizing acoustic noise: enhancing quiet locomotion for quadruped robots in indoor applications. In *Proceedings of 2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Hangzhou, China, October 19–25, 2025, pp. 17972–17979.

[18] Gomes MW, Ahlin K. Quiet (nearly collisionless) robotic walking. In *Proceedings of 2015 IEEE International Conference on Robotics and Automation (ICRA)*, Seattle, USA, May 26–30, 2015, pp. 5761–5766.

[19] Zhang Y, Yao Y, Liu S, Niu Y, Lin C, *et al.* QuietPaw: learning quadrupedal locomotion with versatile noise preference alignment. In *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Hangzhou, China, October 19–25, 2025, pp. 12561–12568.

[20] Panichi E, Ding J, Atanassov V, Yang P, Kober J, *et al.* On-the-fly jumping with soft landing: leveraging trajectory optimization and behavior cloning. *IEEE/ASME Trans. Mechatron.* 2025, 30(4):3142–3151.

[21] Zhu M, Gong D, Zhao Y, Chen J, Qi J, *et al.* Compliant force control for robots: a survey. *Mathmatics* 2025, 13(13):2204.

[22] Joe HM, Oh JH. A robust balance-control framework for the terrain-blind bipedal walking of a humanoid robot on unknown and uneven terrain. *Sensors* 2019, 19(19):4194.

[23] Sygulla F, Rixen D. A force-control scheme for biped robots to walk over uneven terrain including partial footholds. *Int. J. Adv. Rob. Syst.* 2020, 17(1):1729881419897472.

[24] Yi J, Zhu Q, Xiong R, Wu J. Walking algorithm of humanoid robot on uneven terrain with terrain estimation. *Int. J. Adv. Rob. Syst.* 2016, 13(1):35.

[25] Guadarrama-Olvera JR, Kajita S, Cheng G. Preemptive foot compliance to lower impact during biped robot walking over unknown terrain. *IEEE Rob. Autom. Lett.* 2022, 7(3):8006–8011.

[26] Ferrari P, Rossini L, Ruscelli F, Laurenzi A, Oriolo G, *et al.* Multi-contact planning and control for humanoid robots: design and validation of a complete framework. *Rob. Auton. Syst.* 2023, 166:104448.

[27] Nikolic M, Borovac B, Rakovic M. Simulation of foot-ground contact for bipedal robots. In *Proceedings of 2014 23rd International Conference on Robotics in Alpe-Adria-Danube Region (RAAD)*, Smolenice, Slovakia, September 3–5, 2014, pp. 1–6.

[28] Gonzalez P, Buigues G, Mazon AJ. Noise in electric motors: a comprehensive review. *Energies* 2023, 16(14):5311.

[29] Ben Q, Jia F, Zeng J, Dong J, Lin D, *et al.* HOMIE: humanoid loco-manipulation with isomorphic exoskeleton cockpit. *arXiv* 2025, arXiv:2502.13013.