

KiRAS: Keyframe Guided Self-Imitation for Robust and Adaptive Skill Learning in Quadruped Robots

Xiaoyi Wei¹, Peng Zhai^{1,*}, Jiaxin Tu¹, Yueqi Zhang¹,
Yuqi Li¹, Zonghao Zhang¹, Hu Zhou², and Lihua Zhang^{1,*}
Project Website: <https://lainweigh.github.io/kiras>

Abstract—With advances in reinforcement learning and imitation learning, quadruped robots can acquire diverse skills within a single policy by imitating multiple skill-specific datasets. However, the lack of datasets on complex terrains limits the ability of such multi-skill policies to generalize effectively in unstructured environments. Inspired by animation, we adopt keyframes as minimal and universal skill representations, relaxing dataset constraints and enabling the integration of terrain adaptability with skill diversity. We propose Keyframe Guided Self-Imitation for Robust and Adaptive Skill Learning (KiRAS), an end-to-end framework for acquiring and transitioning between diverse skill primitives on complex terrains. KiRAS first learns diverse skills on flat terrain through keyframe-guided self-imitation, eliminating the need for expert datasets; then continues training the same policy network on rough terrains to enhance robustness. To eliminate catastrophic forgetting, a proficiency-based Skill Initialization Technique is introduced. Experiments on Solo-8 and Unitree Go1 robots show that KiRAS enables robust skill acquisition and smooth transitions across challenging terrains. This framework demonstrates its potential as a lightweight platform for multi-skill generation and dataset collection. It further enables flexible skill transitions that enhance locomotion on challenging terrains.

I. INTRODUCTION

Recent advances in deep reinforcement learning (DRL) have led to breakthroughs in quadruped locomotion, enabling quadruped robots to traverse unstructured terrains, such as stairs and obstacles, using proprioception [1], [2]. However, these policies often learn a single terrain-specific solution, overlooking alternative behaviors required by different tasks. This limits their ability to meet diverse demands, such as crawling for concealment during wildlife photography (Fig. 1(b)) or adopting a bipedal stance for inspection tasks (Fig. 1(d)) [3]. This gap underscores the need to integrate **terrain adaptability** with **skill diversity**.

Current approaches to diverse skill learning mainly rely on reward engineering [4] or dataset-driven imitation learning, such as Adversarial Motion Prior (AMP) [5]. Both achieve reliable execution and stable switching of multiple skills on flat terrain, but extending them to complex terrains remains challenging. In unstructured settings, reward engineering must balance skill acquisition with terrain adaptation, greatly increasing design complexity. Imitation-based methods reduce this burden using expert datasets, yet these datasets




Fig. 1: Deployment of KiRAS on Solo-8 and Unitree Go1 robots across diverse environments. (a) Solo-8 robot traverses various obstacles by flexibly switching learned skills. (b) Solo-8 robot exploits crawl for concealment and stilt for expanded field of view. (c) In outdoor unstructured environments, Solo-8 robot demonstrates strong robustness. (d) Unitree Go1 robot acquires bipedal skills to climb a 10 cm step.

strictly constrain **skill representation** and are mostly collected on flat terrain, rarely include rough-terrain scenarios. Consequently, directly applying flat-terrain datasets to rough terrains often leads to overfitting, as terrain-specific motion patterns are poorly captured, and desired motions cannot be reliably reproduced. Although recent studies explore skill-conditioned traversal [6], [7], [8], they still face complex reward design and poor scalability. Thus, achieving terrain adaptability through motion imitation in challenging environments remains an open problem.

Given the scarcity of datasets on complex terrains, a critical question arises: is there a compact, general skill representation that can describe skills with minimal resources, thereby weakening the strong motion constraints imposed by full expert datasets? **Keyframe** from character animation offers one such idea: animators specify a few pivotal poses, and the system interpolates transitions between them, eliminating the need to animate every frame manually [9]. The same idea naturally extends beyond animation: when acquiring new skills, humans also rely on a small set of decisive movements from demonstrations (videos, texts, or mentors), iteratively imitating, refining the successful instances, and generalizing to new situations [10]. Motivated by the shared principle, we treat keyframes as high-level objectives and introduce self-imitation learning (SIL, [11]) as the underlying framework. During training, SIL selects the robot’s high-return trajectories that meet keyframe targets and reinforces them, allowing

¹College of Intelligent Robotics and Advanced Manufacturing, Fudan University, China. (email: weixy23@m.fudan.edu.cn; {pzhai, lihuaazhang}@fudan.edu.cn.)

²Power China Huadong Engineering Corporation Limited, China.

* Corresponding Author.

natural and generalizable motion policies without constraints from fixed demonstrations or large datasets. This enables the robot to better adapt to diverse environmental dynamics and offers a promising path toward unifying terrain adaptability and skill diversity in legged locomotion.

We present **KiRAS** (Keyframe Guided Self-Imitation for Robust and Adaptive Skill Learning), an end-to-end framework for quadruped robots to traverse terrains with diverse skill primitives. Guided by high-level keyframes, KiRAS first acquires diverse skills and transitions via SIL on flat terrain, and then transfers the policy to rough terrains through same policy network for robustness. We further introduce a proficiency-based Skill Initialization Technique that prioritizes under-trained skills to mitigate catastrophic forgetting. Experiments on Solo-8 [12] and Unitree Go1 robots (Fig. 1) demonstrate the method’s ability to perform multi-skill terrain traversal and generalization across platforms.

The main contributions of this work are as follows:

- A keyframe-guided pipeline that extracts high-quality trajectories for SIL to acquire skill primitives, and subsequently trains the same policy network for complex terrains using an Environmental Context Estimator (ECE) conditioned on proprioception.
- A proficiency-driven Skill Initialization Technique that adaptively oversamples under-mastered skills while preserving already acquired behaviors to balance learning.
- Comprehensive evaluations on Solo-8 and Unitree Go1 robots in both simulation and hardware, demonstrating state-of-the-art performance and flexible, robust multi-skill terrain traversal.

II. RELATED WORK

A. Robust DRL Controller without External Sensors

Since the introduction of Rapid Motor Adaptation [13], which demonstrated that DRL policies outperform traditional controllers in robustness for quadrupedal locomotion, increasing attention has been devoted to enabling proprioceptive control and environmental awareness without external sensors. Ji et al. [14] propose a state estimation network to infer linear velocity, foot height, and contact probability directly from onboard observations, reducing reliance on external inputs. Building on this, Nahrendra et al. [2] introduce DreamWaQ, enhancing estimation accuracy with a context-aided estimator. Long et al. [15] develop a model combining regression and contrastive learning to predict velocity and latent states from proprioceptive history. These proprioception-based modules enable the robot to estimate physical quantities not accurately observable through an IMU and motors, improving state awareness, terrain adaptability, and controller robustness.

However, most of these approaches adopt an all-in-one policy that learns a single fixed approach for each terrain via environmental adaptation. While this improves robustness to some extent, it limits the ability to select and switch skills according to task requirements. Consequently, robots rely on a single behavioral pattern per terrain, restricting flexibility

and performance in multi-task scenarios, especially when skill switching is critical.

B. Multi-skill Learning for Legged Robots

Multi-skill learning aims to develop a unified policy network that can acquire and manage multiple skills. In legged robots, Generative Adversarial Imitation Learning (GAIL) [16] is a mainstream approach that reproduces expert state-action distributions through adversarial training, eliminating the need for explicit reward design. AMP [5] extends this idea by enabling skill learning from expert datasets. Building on this foundation, Li et al. propose a series of methods: WASABI [17] infers rewards from rough partial demonstrations; Cassi [18] integrates GAIL with unsupervised skill discovery to isolate individual skills from unlabelled data, while FLD [19] constructs a latent space to improve interpolation and generalization from sparse trajectories. Moreover, Vollenweider et al. [20] employ parallel discriminators for each skill to enable skill transition. Zargarbashi et al. [21] propose a keyframe-guided approach for natural gait generation with temporal keyframe constraints, though it still relies on AMP to compute style rewards.

However, most methods focus on flat terrain and overlook the advantages of diverse skills in complex environments, leaving skill transition for high-level tasks underexplored. While recent studies have begun exploring this direction, they still face limitations. Margolis et al. [4] introduce a multi-skill policy modulated by behavior parameters, but it requires carefully crafted rewards and lacks terrain adaptability. Han et al. [6] and Zhang et al. [7] train multi-gait base networks on large-scale datasets and use residual networks for high-level control. While effective in complex terrains, their multi-stage distillation increases architectural complexity and reduces policy scalability. Huang et al. [8] use a Mixture-of-Experts framework for quadrupedal–bipedal transitions, but skill and task diversity remain limited.

In contrast, KiRAS acquires diverse skills from keyframes without relying on high-quality expert datasets or elaborate reward engineering. It employs an end-to-end training pipeline for a single policy network, balancing performance with architectural simplicity and avoiding the information loss typically introduced by distillation.

III. METHOD

In this section, we introduce KiRAS framework, as illustrated in Fig. 2. We first present the training pipeline, which performs **skill learning** followed by **terrain finetuning**. Then, we describe the Skill Initialization Technique. Finally, we provide some key training details.

A. End-to-end Adaptive Multi-skill Learning Framework

KiRAS enables quadruped robots to perform various skills across diverse terrains, with smooth transitions between them. During **skill learning**, KiRAS learns these skills on flat terrain using keyframe-based SIL, where each skill is represented by *one* single keyframe and encoded into the observation space via a one-hot vector c_i^g . While efficient, this

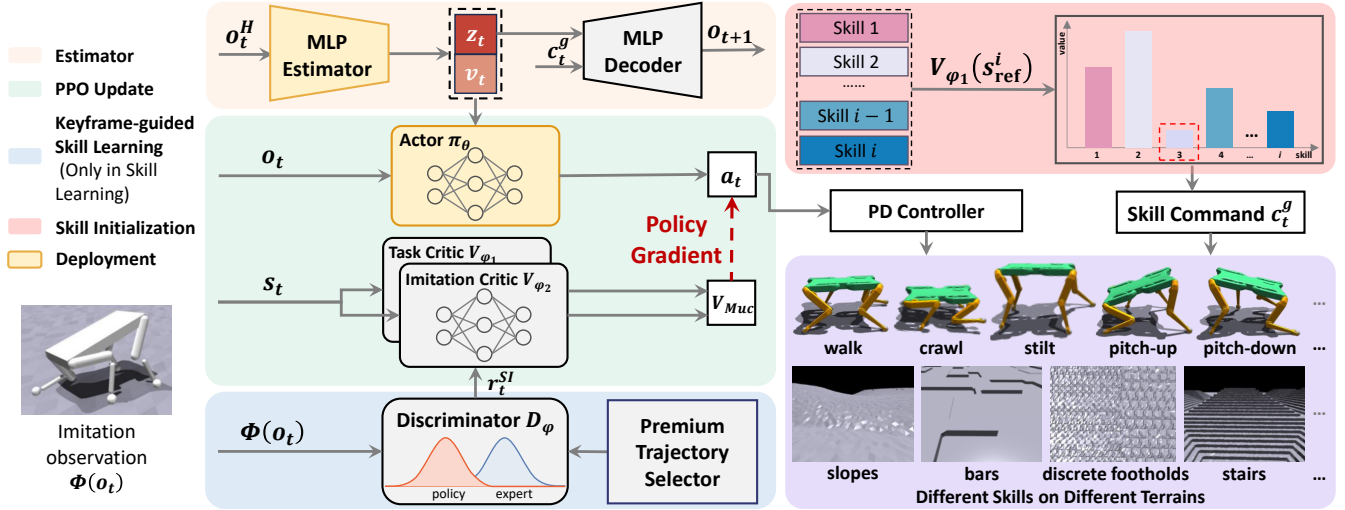


Fig. 2: Training pipeline. KiRAS consists of an end-to-end adaptive multi-skill learning framework and a skill initialization module. The former includes a keyframe-guided skill learning module (used only in skill learning stage), a proprioception estimator, and a PPO actor-critic architecture. The policy outputs joint position targets, which are converted to torques via a PD controller. The Skill Initialization Technique is employed at environment reset to prevent overfitting to simpler skills. The purple box illustrates different skills being trained on different terrains. During deployment, only the networks in the yellow boxes are used, and skill switching is controlled via joystick input.

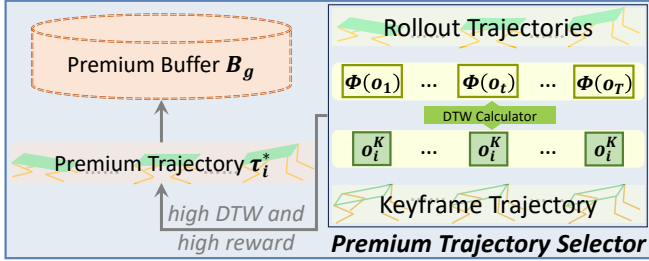


Fig. 3: Details of Premium Trajectory Selector.

minimalist representation limits adaptability: joint positions in a keyframe remain fixed, failing to capture differences in execution between, for example, walking on flat ground and ascending stairs. To overcome this limitation, the subsequent **terrain finetuning** disables SIL module and continues training the same policy network to acquire robustness and flexibility for terrain traversal. Unlike conventional two-stage frameworks that rely on distillation, KiRAS achieves joint learning of skills and terrain adaptability through the keyframe mechanism, mitigating inefficiency and information loss observed in traditional pipelines.

1) *Keyframe-guided Skill Learning*: In skill learning, we use keyframes as inputs for SIL to facilitate adversarial skill acquisition. Our approach builds on the WASABI [17] framework. Keyframes can be seen as special cases of partial demonstration containing a single timestep. However, only repeating this static frame to form transitions can cause overfitting to the posture and lead to temporally incoherent actions. To solve this, we use premium trajectories collected through exploration as demonstrations for SIL.

Fig. 3 illustrates the Premium Trajectory Selector. For each skill i , we maintain a premium buffer $B_g = \{\tau_i^*\}$, where τ_i^* denotes premium trajectories. Premium buffer is initialized with trajectories generated by repeating the keyframe K of skill i for T timesteps, referred to as *keyframe trajectories*: $\{o_i^K\}^T$, where o_i^K contains partial observations such as the

root state and joint positions. Inspired by [22], we select a *premium trajectory* $\tau_\pi = (\Phi(o_1), \Phi(o_2), \dots, \Phi(o_T))$ for skill i based on two criteria: (i) the combined reward r_t^c accumulated over the rollout (detailed in Section. III-A.2), measuring performance and alignment with desired behaviors, and (ii) the Dynamic Time Warping (DTW) distance to the keyframe trajectory, which measures structural similarity in the imitation space. Here, Φ denotes an operator that extracts partial information from the full observations into the imitation space, matching the observation modality of the keyframe. The combined evaluation is defined as:

$$J(\tau_\pi) = \sum_{t=0}^T r_t^c + d^{\text{DTW}}(\tau_\pi, \{o_i^K\}^T). \quad (1)$$

After each episode, if $J(\tau_\pi)$ exceeds the current maximum score ϵ_i of the premium trajectory for skill i in the premium buffer, the trajectory is added to B_g ; otherwise it is discarded. Keyframe trajectories are consistently retained in the premium buffer as references.

We then train a discriminator D_ψ to distinguish state transitions $(\Phi(o_{t-1}), \Phi(o_t))$ from the premium buffer distribution B_g or the policy distribution d_π . Following the same formulation as in WASABI, SIL reward r_t^{SI} encourages the policy to generate transitions resembling high-return trajectories, guiding self-imitation and effective skill learning.

2) *Rewards and Multi-critic Architecture*: In KiRAS, the total reward comprises 4 components: SIL reward r_t^{SI} for learning skill primitives, the termination penalty r_t^T , the regularization reward r_t^R for stable real-world deployment, and a novel residual penalty r_t^{res} that constrains joint outputs. More details about r_t^T and r_t^R can be found in [17], and the residual term r_t^{res} is defined as:

$$r_t^{res} = \begin{cases} 0, & 0 \leq t < T_1, \\ -\|a - a_{\text{flat}}\|, & T_1 \leq t < T_2, \end{cases} \quad (2)$$

where T_1 and T_2 denote the timesteps of skill learning and terrain finetuning, respectively, and a_{flat} is the action produced in skill learning under identical observations. The reward aligns actions in terrain finetuning with those from skill learning to prevent detrimental deviations.

During terrain finetuning, all terms except r_t^{SI} remain active to guide robust terrain adaptation. Relying on a single critic to estimate a composite value function can overload by heterogeneous reward signals, causing high variance in value estimates, and unstable policy gradients. Moreover, residual gradients from r_t^{SI} can persist and disrupt terrain learning.

To address these issues, we adopt a multi-critic (MuC) architecture inspired by [21], where each critic V_{ϕ_i} is responsible for estimating a specific subset of rewards. Specifically, the **Task Critic** V_{ϕ_1} estimates the cumulative return from the combined reward $r_t^c = \omega^T r_t^T + \omega^R r_t^R + \omega^{res} r_t^{res}$, while the **Imitation Critic** V_{ϕ_2} estimates the cumulative return from the imitation reward $\omega^{SI} r_t^{SI}$. Each critic is trained independently to specialize in its corresponding reward signal. At each update, we compute normalized advantages \hat{A}_i with mean μ_i and standard deviation σ_i , and combine them as $\hat{A}_{MuC} = \sum_{i=1}^2 \omega_i \frac{\hat{A}_i - \mu_i}{\sigma_i}$. However, unlike prior work, we dynamically adjust the mixing weights ω_1, ω_2 over training iterations t with hyperparameter σ :

$$\omega_1 = \begin{cases} \frac{\sigma}{T_1} t + (1 - \sigma)t, & 0 \leq t < T_1, \\ 1, & T_1 \leq t < T_2, \end{cases} \quad (3)$$

$$\omega_2 = 1 - \omega_1,$$

where ω_1 and ω_2 denote the task advantage and imitation advantage weight. This linear schedule shifts emphasis from skill imitation to terrain adaptation as training progresses.

3) *Environment Context Estimator*: To enhance terrain robustness and adaptability, we introduce ECE that enriches the policy network with proprioceptive context. ECE consumes the past $H = 5$ steps of proprioceptive observations o_t^H to predict the body velocity v_t and produce a latent vector z_t , which is used to reconstruct the next observation o_{t+1} . Previous studies have demonstrated that accurate velocity estimation and implicit environmental representation are critical for locomotion [2], [15]. However, different skills correspond to distinct postures and contact dynamics, and conditioning on the skill command c_t^g helps the estimator capture these variations across skills.

To address this, we employ a Conditional Variational Autoencoder (CVAE) that incorporates the skill command as a context prior. The observation reconstruction is defined as $\hat{o}_{t+1} = \text{Decoder}(z_t, c_t^g)$, and the velocity prediction is supervised using MSE. The combined loss is

$$\mathcal{L}_{ECE} = \text{MSE}(v_t, \hat{v}_t) + \text{MSE}(o_{t+1}, \hat{o}_{t+1}) + \beta \text{DKL}(q(z_t | o_t^H, c_t^g) \| p(z_t | c_t^g)), \quad (4)$$

where \hat{v}_t and \hat{o}_{t+1} are the predicted velocity and observation, respectively. $q(z_t | o_t^H, c_t^g)$ is the posterior distribution and $p(z_t | c_t^g)$ is a skill-conditioned Gaussian prior.

Since skill learning is performed on flat terrain and terrain finetuning on complex terrain, early ECE convergence on flat

terrain may hinder the modeling of uneven environments. Thus, we adopt adaptive sampling (AdaBoot), which conditionally updates the estimator based on reward ratio [2], improving generalization to complex terrains.

B. Skill Initialization Technique

One of the key challenges in multi-skill learning is exploration bias, where the policy tends to overfit to easier sub-skills while neglecting those that are more difficult to acquire [23]. To counteract this bias, we introduce the Skill Initialization Technique that probabilistically favors undertrained skills. Taking advantage of MuC framework, we leverage Task Critic to evaluate the performance of each skill and dynamically adjust their sampling probabilities based on their task values, which enables balanced training across all skills within the policy network, mitigating catastrophic forgetting and promoting comprehensive skill development.

Concretely, at the start of an episode, we associate each skill i with a reference state s_i^{ref} , from which we compute its task value $V_{\phi_1}(s_i^{\text{ref}})$ using the Task Critic. In practice, we set s_i^{ref} to the skill's corresponding keyframe s_i^K . The probability of selecting skill i at the beginning of an episode is:

$$p(\text{Skill} = i) = \frac{1}{N-1} \left(1 - \frac{V_{\phi_1}(s_i^{\text{ref}})}{\sum_{m=1}^N V_{\phi_1}(s_m^{\text{ref}})} \right), \quad (5)$$

where N denotes the total number of skills. Skills with lower task values are more likely to be sampled, encouraging focused improvement. Using relative performance, the policy can adaptively shift attention toward weaker skills.

To ensure all skills remain represented during training, we perform a coverage check. If any skill has not been sampled recently, we substitute one of the most frequently sampled skill to maintain balanced skill distribution. Thanks to this selector-based mechanism, multiple skills can be learned simultaneously within a single policy.

C. Training Details

1) *Terrain Finetuning*: We design 5 terrain types: flat, slopes, bars, discrete footholds, and stairs, where the latter four are rough terrains. The discrete footholds simulate scenarios where the robot must recover from leg entrapment in gaps. Each type includes 10 difficulty levels, with complexity progressively increased using the terrain curriculum proposed in [24]. All terrain types are generated before training, with the robot selecting the corresponding terrain during learning. To encourage the robot to lift legs over higher obstacles, we add fractal noise to rough terrains and embed the curriculum with a noise amplitude range of $[0.02, 0.07]$.

2) *State Initialization*: The design of the initial state plays a crucial role in guiding the robot into more informative regions, accelerating convergence while balancing exploration and exploitation. To leverage keyframes, we randomly initialize the robot's joint positions with the keyframe s_i^K of the selected skill i (Section. III-B). Notably, the initialized state can differ from the state of the target skill described in Section III-B, which further facilitates seamless skill transitions and mitigates forgetting.

IV. EXPERIMENTS

The following experiments aim to address:

- Does each component of KiRAS enhance skill diversity or policy robustness? Can KiRAS match the performance of state-of-the-art open-source algorithms under the same conditions? (In Section. IV-B)
- Can KiRAS learn multiple skills simultaneously, with each achieving ideal performance? (In Section. IV-C)
- Is KiRAS capable of traversing complex indoor and outdoor environments with stable and robust performance? How does flexible skill switching benefit locomotion in such environments? (In Section. IV-D)
- How well does KiRAS generalize? Can it transfer to quadruped platforms with different DOFs and support a broader range of skill types? (In Section. IV-E)

A. Experimental Setup

We use Isaac Gym [25] as the simulation platform. In both simulation and real-world experiments, we deploy the Solo-8 robot to perform 5 skills: walking, stilting, crawling, pitching-up, and pitching-down, with their transitions. Each skill has a specific base height and pitch angle summarized in TABLE I. We train the policy with $T_1 = 6,000$ iterations for skill primitives acquisition and $T_2 = 30,000$ iterations for robust adaptation, using 4,096 parallel environments on an NVIDIA RTX 4090 GPU within 12 hours. After training, we export the policy as an ONNX model (under 500 kB), so it can run on a low-cost onboard computer at 50 Hz.

TABLE I: Parameters of Different Skills.

Skills	Walk	Crawl	Stilt	Pitch-up	Pitch-down
Base Height (cm)	20	10	30	20	20
Pitch Angle (deg)	0	0	0	-15	15

B. Ablation and Comparison Experiments

1) *Analysis of Compared Methods:* To comprehensively evaluate the performance of KiRAS, we design 8 comparative baselines in simulation platform:

- **KiRAS w/o SIL:** The policy excludes the Premium Trajectory Selector, manually designing reward functions for each skill based on base height and pitch angle.
- **KiRAS w/o MuC:** The policy excludes MuC and employs a single critic network for value calculation.
- **KiRAS w/o ECE:** The policy excludes ECE while retaining o_t^H as proprioceptive observation input.
- **KiRAS w/o Skill Initialization:** The policy excludes the Skill Initialization Technique and randomly selects a skill to train at the start of each episode.
- **Cassi [18]:** A method that learns skills from unlabeled expert datasets, which collected from open-source algorithms [17], [18], [26].
- **WTW [4]:** A method encoding a structured family of locomotion policies that solve training tasks in diverse approaches, thereby enabling skill diversity.
- **DreamWaQ [2]:** A VAE-based method for estimating velocity and future observations, modified with manually designed reward functions to distinguish skills.

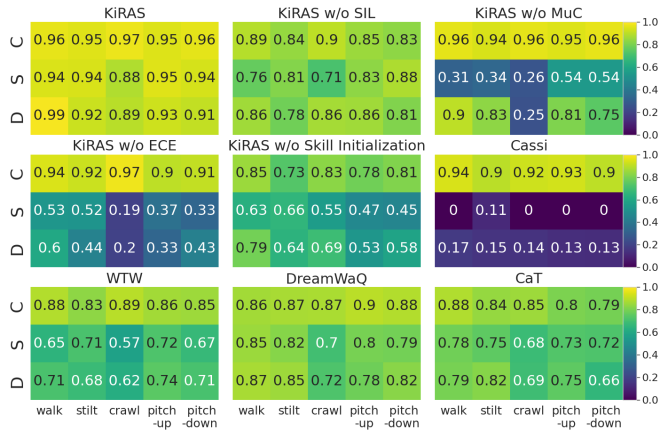


Fig. 4: Heatmap of ablation and comparison experiment results. The vertical axis "C" denotes the cosine similarity between each skill trajectory and its corresponding keyframe, with higher values indicating greater similarity. "S" and "D" represent the success rates of traversing steps and discrete footholds, respectively, where higher values indicate more robust policies.

- **CaT [27]:** CaT employs an early-terminated DRL algorithm for stable exploration, modified with manually designed constraints to distinguish skills.

All algorithms are trained for 36,000 iterations with identical terrain settings (depicted in Section. III-C) and rewards unless otherwise specified. After training, we compute the cosine similarity (C) between each skill and its corresponding keyframe, which measures the dot-product similarity of trajectories and serves as an indicator of skill acquisition. We further evaluate performance over 1,000 trials across arbitrary speeds and orientations, measuring success rates on steps (S) and discrete footholds (D), which quantify terrain adaptability. The resulting heatmap is presented in Fig. 4.




Fig. 5: Comparative results. (a) KiRAS w/o Skill Initialization learns 5 skills (walk, crawl, stilt, pitch-up, pitch-down) with little distinction. (b) DreamWaQ's crawl and pitch-down use knee-ground contact (red circles). (c) KiRAS w/o MuC fails to lift legs over obstacles.

With an effective training pipeline and model architecture, KiRAS achieves the highest imitation similarity and superior terrain performance. As shown in the first row of each subfigure, KiRAS maintains dataset-level imitation quality comparable to Cassi, but without relying on large expert datasets. In contrast, Cassi suffers extremely low traversal success due to being overly constrained by the dataset and lacking terrain awareness. This also partially affects Cassi's skill imitation, leading the policy to converge to unnatural behaviors on challenging terrains.

Removing the Skill Initialization Technique collapses skills into nearly identical gaits that diverge from the keyframes, as illustrated in Fig. 5(a), highlighting its role in preventing forgetting. Both KiRAS w/o SIL and DreamWaQ

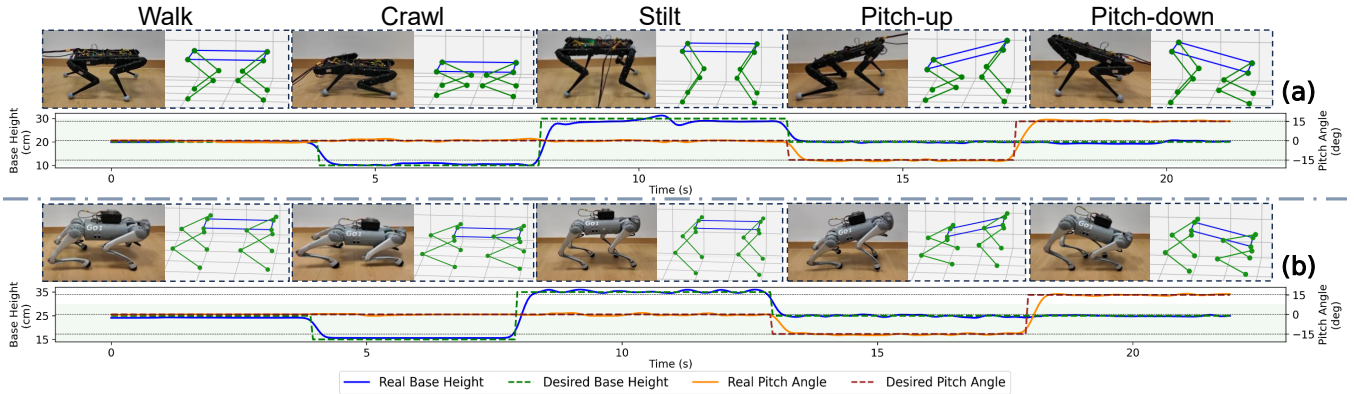


Fig. 6: Real-world experiments validating skill acquisition with KiRAS on the Solo-8 (a) and Unitree Go1 (b) robots. Each black dashed box shows skill execution on hardware, with the left image showing the real robot and the right showing the keyframe used for training. In all experiments, robots are oriented to the right as the forward direction.

constrain skills via rewards. While effective, they over-constrain the base height or orientation, resulting in the robot contacting the terrain with its knees rather than its feet, especially when the base height is relatively low, as shown in Fig. 5(b). This slightly improves success rates but severely compromises hardware safety. Without MuC, the policy remains stuck in static keyframes, failing to learn dynamic foot-lifting and thus struggling on steps, as shown in Fig. 5(c). Without ECE, the policy loses proprioceptive terrain awareness and performs like WTW; although WTW distinguishes skills, neither approach has terrain perception and both fail on challenging terrains. CaT relies on constrained terminations but is not designed for traversal, limiting its performance.

2) *Effectiveness of Skill Initialization Technique*: This experiment evaluates the Skill Initialization Technique for balancing multi-skill learning. Fig. 7 shows the evolution of sampling probability $p(\text{Skill} = i)$ for 5 skills over 36,000 iterations. During skill learning, all skills converge at about 4,500 iterations with probabilities near 0.2. During terrain finetuning, variations demand relearning adaptability. Among these skills, walk and crawl converge faster due to lower center of mass, reduced motion range, and balance demands, leading to lower probabilities. In contrast, stilt requires higher center of mass and extended limbs, demanding precise control to avoid instability and thus receiving higher probability. Pitch-up and pitch-down require stable pitch and altered force distribution, increasing correction difficulty and yielding the highest probabilities. The Solo-8’s front-heavy design makes pitch-down more prone to slipping or overload, raising its probability. To prevent forgetting, all skills maintain non-zero sampling throughout training. These results show KiRAS adaptively adjusts sampling to skill difficulty, ensuring balanced multi-skill learning.

C. Skill Imitation Sim-to-real Experiments

To evaluate skill acquisition, we extract the policy and deploy it on the real Solo-8 robot, comparing the skills with the corresponding keyframes. Moreover, we record the base height and pitch angle while the robot executes skills and transitions. As shown in Fig. 6(a), the executed




Fig. 7: Sampling probability evolution for each skill during training, computed using 4 random seeds. Curves and shaded areas represent the mean and standard deviation across seeds, respectively.

motions closely match the designed keyframes, confirming the accuracy in reproducing the intended motion style. The tracked values further align with TABLE I, verifying precise behaviors under keyframe constraints. In addition, during transitions between different skills, the trajectories remain smooth and continuous, without noticeable oscillations or delays, demonstrating remarkable stability and controllability in dynamic switching.

Overall, these results show that KiRAS efficiently learns and transitions among multiple skills from keyframe guidance alone. The comparable similarity to large-scale datasets highlights the potential as a compact, keyframe-driven platform for generating multi-skill quadruped demonstrations, enabling efficient collection of complex locomotion tasks.

D. Indoor and Outdoor Sim-to-real Experiments

Extensive indoor and outdoor experiments demonstrate KiRAS’s effectiveness, robustness, and versatility in sim-to-real transfer, enabling traversal of diverse terrains with multiple skills. These results highlight the importance of flexible skill switching for locomotion in complex environments. To our knowledge, this is the first work to achieve agile locomotion control for an 8-DoF quadruped robot on complex terrain.

1) *Locomotion Ability Test*: Fig. 8 shows the Solo-8 performing multi-skill locomotion: (a) walking across wedge-shaped obstacles and soft ground; (b) climbing a 10 cm step with pitch-up; (c) crawling through a 15 cm-high stone-filled passage; and (d) crawling through pit-filled grass and a bicy-




Fig. 8: Solo-8 robot traversing diverse and challenging terrains in indoor and outdoor environments with different skills.

cle obstacle, while descending steps and traversing slippery surfaces by the walk skill. All tasks require lifting the legs to overcome elevated obstacles, which emerges in the terrain finetuning even though it is not explicitly encoded in the keyframes. Solo-8 is compact (2 kg), with limited joint torque (1 Nm, about 1/30 of Unitree Go1) and no hip joints, making rough-terrain locomotion challenging. However, KiRAS enables it to complete demanding tasks (Fig. 8(a,b)). The lack of joint limits further allows extremely low postures, enabling traversal through narrow gaps (Fig. 8(c,d-upper); Fig. 1(c)). Moreover, its lightweight design provides stability on slippery surfaces through improved center-of-mass control (Fig. 8(d-lower)). These results demonstrate the robustness of the policies trained with KiRAS, enabling the robot to traverse complex terrains with diverse skills.

2) *Terrain Adaptation of Individual Skills*: To verify that specific skills better adapt to challenging terrains, we conduct 50 experiments on a real Solo-8 robot, using different skills to climb 45° slopes and traverse 15 cm steps. These tasks are highly challenging for the Solo-8 robot. The success rates, summarized in TABLE II, show that pitch-down achieves the highest rate on steep slopes by lowering the front body to reduce tipping and extending the hind legs for stronger propulsion. In contrast, pitch-up excels on stairs, as lifting the forelegs facilitates step clearance while stabilizing the center of mass. Other skills exhibit limitations in either stability or obstacle clearance, resulting in lower success rates.

TABLE II: Success rates of Solo-8 robot traversing terrains with different skills. The highest rates are in bold.

Skills	Walk	Crawl	Stilt	Pitch-up	Pitch-down
Slope	0.84	0.68	0.80	0.70	0.9
Upstairs	0.80	0.54	0.44	0.88	0.58
Downstairs	0.88	0.46	0.70	0.94	0.56

To further illustrate this, Fig. 1(a) shows the robot employing different skills: crawl for a low bar, stilt for a tall barrier, pitch-down for ascending a slope, and pitch-up for descending stairs. The first two tasks are posture-constrained, with crawl and stilt required by geometry, demonstrating the role of base height in restricted settings. The latter highlight



Fig. 9: (a) Transition from the walk skill to the biped skill. (b) Keyframe for the biped skill used for imitation. (c) Snapshots of executing the biped skill to traverse various indoor and outdoor terrains.

flexibility, consistent with quantitative results above. These results confirm that skills provide task-specific solutions to improve stability across diverse terrains.

E. Generalization Experiments

1) *Generalization to Robot with Different DOFs*: Since KiRAS relies solely on single-frame keyframes, its reference data is easily accessible, enabling cross-platform skill training. We train a policy on Unitree Go1 robot by adjusting joint positions of Solo-8 keyframes, achieving accurate multi-skill reproduction (Fig. 6(b)). Because Unitree Go1 is larger than Solo-8, all keyframes are designed with a 5 cm higher base. Similar to Section IV-C, we record base height and pitch angle during continuous trajectories. Results show that KiRAS can preserve precise skill execution and smooth transitions on a robot with different DOFs.

2) *Generalization to New Skills by New Keyframes*: To evaluate the generalization of KiRAS, we train an additional skill that can be represented by one keyframe: biped (Fig. 9(b)). This task requires the robot to stand upright on its hind legs, which is more challenging due to reduced stability, requiring continuous gait adjustments for balance. In our experiments, we extend the existing policy trained on 5 skills by incorporating the keyframes of the new skill, followed by

training for $T_1 = 10,000$ and $T_2 = 30,000$ timesteps to progressively integrate it into the original framework.

As shown in Fig. 9(a), when transitioning from walking to bipedal stance, the robot first raises its front body and shifts the center of mass to the hind legs. It then pushes off with its forelegs to gain upward momentum, bringing the base close to vertical. At this stage, the robot relies on continuous hind-leg adjustments to maintain balance. This demonstrates that KiRAS can effectively learn more challenging skills while maintaining stable control.

Furthermore, we evaluate the robot’s biped skill on a variety of unstructured terrains, including smooth slopes, discrete stepping stones, sponges, stairs, sloped grassy surfaces, and playgrounds with pits and gravel, as illustrated in Fig. 1(d), and Fig. 9(c). Among these, stairs and discrete terrains have proven to be the most challenging. Nevertheless, thanks to the robustness of the policy, the bipedal robot can safely traverse stairs up to 10cm in height. These results further demonstrate that KiRAS enables stable locomotion in complex environments.

V. CONCLUSIONS

This work presents KiRAS, a keyframe-driven end-to-end DRL framework that unifies multi-skill learning and terrain adaptability. KiRAS acquires diverse skills from single keyframes within a unified policy, and progressively masters locomotion tasks through skill primitives. The Skill Initialization Technique further mitigates catastrophic forgetting by efficiently selecting subsequent training actions. Experiments show that KiRAS achieves strong skill diversity and terrain adaptability, matching open-source methods built on large datasets. It also generalizes to quadruped platforms with varying DOFs and new skill types, highlighting its potential as a lightweight platform for multi-skill generation and dataset collection. Future work will explore adaptive skill switching using environmental cues, and extend the framework to navigation tasks with exteroceptive sensing.

REFERENCES

- [1] T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, “Learning robust perceptive locomotion for quadrupedal robots in the wild,” *Science Robotics*, vol. 7, no. 62, p. eabk2822, 2022. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.abk2822>
- [2] I. M. A. Nahrendra, B. Yu, and H. Myung, “Dreamwaq: Learning robust quadrupedal locomotion with implicit terrain imagination via deep reinforcement learning,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 5078–5084.
- [3] Y. Niu, Y. Zhang, M. Yu, C. Lin, C. Li, Y. Wang, Y. Yang, W. Yu, T. Zhang, Z. Li, J. Francis, B. Chen, J. Tan, and D. Zhao, “Human2locoman: Learning versatile quadrupedal manipulation with human pretraining,” in *Robotics: Science and Systems (RSS)*, 2025.
- [4] G. B. Margolis and P. Agrawal, “Walk these ways: Tuning robot control for generalization with multiplicity of behavior,” in *Conference on Robot Learning*. PMLR, 2023, pp. 22–31.
- [5] X. B. Peng, Z. Ma, P. Abbeel, S. Levine, and A. Kanazawa, “Amp: adversarial motion priors for stylized physics-based character control,” *ACM Transactions on Graphics*, vol. 40, no. 4, p. 1–20, Jul. 2021. [Online]. Available: <http://dx.doi.org/10.1145/3450626.3459670>
- [6] L. Han, Q. Zhu, J. Sheng, C. Zhang, T. Li, Y. Zhang, H. Zhang, Y. Liu, C. Zhou, R. Zhao *et al.*, “Lifelike agility and play in quadrupedal robots using reinforcement learning and generative pre-trained models,” *Nature Machine Intelligence*, vol. 6, no. 7, pp. 787–798, 2024.
- [7] Z. Zhang, C. Li, T. Miki, and M. Hutter, “Motion priors reimaged: Adapting flat-terrain skills for complex quadruped mobility,” *arXiv preprint arXiv:2505.16084*, 2025.
- [8] R. Huang, S. Zhu, Y. Du, and H. Zhao, “Moe-loco: Mixture of experts for multitask locomotion,” *arXiv preprint arXiv:2503.08564*, 2025.
- [9] D. Sturman, “Interactive key frame animation of 3-d articulated models,” in *Graphics Interface*, vol. 86, 1984.
- [10] A. J. Ijspeert, J. Nakanishi, H. Hoffmann, P. Pastor, and S. Schaal, “Dynamical movement primitives: learning attractor models for motor behaviors,” *Neural computation*, vol. 25, no. 2, pp. 328–373, 2013.
- [11] J. Oh, Y. Guo, S. Singh, and H. Lee, “Self-imitation learning,” 2018. [Online]. Available: <https://arxiv.org/abs/1806.05635>
- [12] F. Grimmering, A. Meduri, M. Khadiv, J. Viereck, M. Wuthrich, M. Naveau, V. Berenz, S. Heim, F. Widmaier, T. Flayols, J. Fiene, A. Badri-Sprowitz, and L. Righetti, “An open torque-controlled modular robot architecture for legged locomotion research,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, p. 3650–3657, Apr. 2020. [Online]. Available: <http://dx.doi.org/10.1109/LRA.2020.2976639>
- [13] A. Kumar, Z. Fu, D. Pathak, and J. Malik, “Rma: Rapid motor adaptation for legged robots,” *arXiv preprint arXiv:2107.04034*, 2021.
- [14] G. Ji, J. Mun, H. Kim, and J. Hwangbo, “Concurrent training of a control policy and a state estimator for dynamic and robust legged locomotion,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4630–4637, 2022.
- [15] J. Long, Z. Wang, Q. Li, J. Gao, L. Cao, and J. Pang, “Hybrid internal model: Learning agile legged locomotion with simulated robot response,” *arXiv preprint arXiv:2312.11460*, 2023.
- [16] J. Ho and S. Ermon, “Generative adversarial imitation learning,” *Advances in neural information processing systems*, vol. 29, 2016.
- [17] C. Li, M. Vlastelica, S. Blaes, J. Frey, F. Grimmering, and G. Martius, “Learning agile skills via adversarial imitation of rough partial demonstrations,” in *Conference on Robot Learning*. PMLR, 2023, pp. 342–352.
- [18] C. Li, S. Blaes, P. Kolev, M. Vlastelica, J. Frey, and G. Martius, “Versatile skill control via self-supervised adversarial imitation of unlabeled mixed motions,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 2944–2950.
- [19] C. Li, E. Stanger-Jones, S. Heim, and S. Kim, “FId: Fourier latent dynamics for structured motion representation and learning,” *arXiv preprint arXiv:2402.13820*, 2024.
- [20] E. Vollenweider, M. Bjelonic, V. Klemm, N. Rudin, J. Lee, and M. Hutter, “Advanced skills through multiple adversarial motion priors in reinforcement learning,” 2022. [Online]. Available: <https://arxiv.org/abs/2203.14912>
- [21] F. Zargarbashi, J. Cheng, D. Kang, R. Sumner, and S. Coros, “Robotkeyframing: Learning locomotion with high-level objectives via mixture of dense and sparse rewards,” *arXiv preprint arXiv:2407.11562*, 2024.
- [22] J. Tu, X. Wei, Y. Zhang, T. Hou, X. Gao, Z. Dong, P. Zhai, and L. Zhang, “Continuous control of diverse skills in quadruped robots without complete expert datasets,” *arXiv preprint arXiv:2503.03476*, 2025.
- [23] M. Riemer, I. Cases, R. Ajemian, M. Liu, I. Rish, Y. Tu, and G. Tesaro, “Learning to learn without forgetting by maximizing transfer and minimizing interference,” *arXiv preprint arXiv:1810.11910*, 2018.
- [24] N. Rudin, D. Hoeller, P. Reist, and M. Hutter, “Learning to walk in minutes using massively parallel deep reinforcement learning,” in *Conference on robot learning*. PMLR, 2022, pp. 91–100.
- [25] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa, and G. State, “Isaac gym: High performance gpu-based physics simulation for robot learning,” 2021. [Online]. Available: <https://arxiv.org/abs/2108.10470>
- [26] Y. Fuchioka, Z. Xie, and M. Van de Panne, “Opt-mimic: Imitation of optimized trajectories for dynamic quadruped behaviors,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 5092–5098.
- [27] E. Chane-Sane, P.-A. Leziart, T. Flayols, O. Stasse, P. Souères, and N. Mansard, “Cat: Constraints as terminations for legged locomotion reinforcement learning,” in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 13 303–13 310.