

Look Forward to Walk Backward: Efficient Terrain Memory for Backward Locomotion with Forward Vision

Shixin Luo¹, Songbo Li¹, Yuan Hao¹, Yaqi Wang², Jun Zheng³, Jun Wu¹, and Qiuguo Zhu^{*1}

Abstract—Legged robots with egocentric forward-facing depth cameras can couple exteroception and proprioception to achieve robust forward agility on complex terrain. When these robots walk backward, the forward-only field of view provides no preview. Purely proprioceptive controllers can remain stable on moderate ground when moving backward but cannot fully exploit the robot’s capabilities on complex terrain and must collide with obstacles. We present Look Forward to Walk Backward (LF2WB), an efficient terrain-memory locomotion framework that uses forward egocentric depth and proprioception to write a compact associative memory during forward motion and to retrieve it for collision-free backward locomotion without rearward vision. The memory backbone employs a delta-rule selective update that softly removes then writes the memory state along the active subspace. Training uses hardware-efficient parallel computation, and deployment runs recurrent, constant-time per-step inference with a constant-size state, making the approach suitable for onboard processors on low-cost robots. Experiments in both simulations and real-world scenarios demonstrate the effectiveness of our method, improving backward agility across complex terrains under limited sensing.

I. INTRODUCTION

Legged robots with egocentric forward-facing depth cameras can tightly couple exteroception and proprioception to realize robust forward agility, including jumping, climbing, and negotiating tight spaces [1], [2], [3], [4]. When the robot must walk backward, the forward-only field of view provides no preview. Purely proprioceptive controllers can remain stable on moderate terrain, but they follow a feel-first, adapt-later pattern that often collides with obstacles and underperforms in complex scenes where anticipating contact is essential [5]. This raises a concrete question: after traversing a scene forward once, can the robot equipped only with a forward-facing camera reuse a compact terrain memory to anticipate obstacles and achieve collision-free backward locomotion without any rearward exteroception?

One seemingly straightforward option is to insert explicit hierarchical mapping, for example elevation or voxel maps, combined with SLAM for long-term memory for backward locomotion [6]. However, in highly dynamic locomotion,

This work was supported by the National Key R&D Program of China (Grant No. 2022YFB4701502), the “Leading Goose” R&D Program of Zhejiang (Grant No. 2023C01177), and the 2035 Key Technological Innovation Program of Ningbo City (Grant No. 2024Z300).

¹The authors are with Institute of Cyber-Systems and Control, Zhejiang University, 310027, China.

²The author is with College of Artificial Intelligence, Nankai University, 300350, China.

³The author is with Hangzhou Public Library, 310016, China.

*Qiuguo Zhu (qgzhu@zju.edu.cn) is the corresponding author.

The supplementary video is available at <https://youtu.be/qWlKtK1telo>.




Fig. 1: We deploy our policy in real-world environments, demonstrating agile forward and backward motion over discrete terrains that combine steps and gaps. The robot memorizes the surrounding terrain during forward motion and retrieves useful information from the memory state when moving backward, enabling highly dynamic backward locomotion using only a forward-facing depth camera.

such approaches add latency, accumulate drift, and propagate integration noise. This motivates a single-stage training paradigm that learns to store and retrieve terrain memory directly within the control loop, reduces the representational mismatch between modules, and uses a lightweight backbone that fits onboard processors on low-cost robots. Prior work has shown the potential of one-stage pipelines incorporating an asymmetric actor-critic and a LSTM/GRU-based estimator in blind and perceptive locomotion tasks [7], [5], [8].

Classical recurrent units such as LSTM and GRU have therefore been popular in resource-constrained robot learning. They offer constant per-step inference time with respect to context length, a fixed-size recurrent state, and the ability to span long episodes by passing the fixed-size memory hidden state between rollouts during on-policy data collection while truncating gradients at rollout boundaries for Truncated Backpropagation Through Time (TBPTT) [9] to reduce computational cost and memory usage. However, stacked LSTMs/GRUs often suffer from vanishing or exploding gradients, weak long-range modeling, and poor scaling at large batch or sequence lengths, which limits their utility when long and cluttered traversals require high-capacity memory [10], [11].

By contrast, decoder-only self-attention Transformers excel at long-range dependencies and scale incredibly well [12]. Several works on locomotion have adopted Transformer backbones, but the effective memory horizons remain too short to support long-distance backward locomotion that relies on stored terrain [13], [14], [15]. The bottleneck is practical: training with backbones like Transformer-XL requires storing KV states to preserve memory from previous

rollout segments [16], [17], so GPU memory grows with the desired context length and forces short memory windows [18]. This is a disadvantage relative to RNNs, which carry only a constant-size recurrent state. At inference time, attention with KV caching incurs per-step cost that increases with context length since the model attends over all cached keys, which conflicts with the need for stable, constant-time control cycles on real robots and again encourages aggressive memory context window clipping. Moreover, because memory is represented per time step, every step appends KV entries. In our task, if the robot waits in place, the local terrain map should remain unchanged, but step-wise caching keeps growing with redundant entries until the cache limit is reached, at which point earlier terrain is forgotten simply due to limited KV cache capacity.

Modern recurrent backbones have recently closed much of the gap between classic gated RNNs and self-attention Transformers [19], [20], [21], [22]. Unlike stacked LSTMs/GRUs, they train efficiently with parallel or chunkwise parallel algorithms and exhibit better stability and scaling at longer horizons. Unlike decoder-only self-attention Transformers, they require no KV cache, so inference time per step and memory footprint remain essentially constant with respect to context length, which keeps control cycles stable for onboard processors. These backbones also strengthen long-range modeling through data-aware state updates and structured transitions, yielding compact, durable associative memory that is particularly suitable for storing and retrieving traversed terrain during backward locomotion.

From the perspective of memory mechanisms developed for language models, our setting can be viewed as an in-context retrieval task: forward motion must write terrain cues that remain retrievable much later when the robot reverses. Plain linear attention is not suitable because it only accumulates key-value associations and lacks a principled removal mechanism, which leads to overload and degraded retrieval as sequences grow. A common fix is to add an explicit forget gate. In language models, a forget gate is valuable for rapidly clearing outdated or irrelevant information during context switches. In backward locomotion with forward-only vision, however, the memory should behave more like an implicit SLAM store, where information encountered long ago can become critical when the robot moves back, and aggressive decay risks erasing exactly that information. Training with rollout windows further amplifies this effect because state is detached across windows, so gradient does not propagate globally and explicit forgetting accelerates erosion of distant but relevant terrain memory. Among modern recurrent models, delta-rule variants such as DeltaNet perform especially well on in-context retrieval because each step first removes outdated content tied to the current input and then adds a calibrated update from the new observation [23], [24], [25]. This selective overwrite avoids uncontrolled accumulation without invoking rapid global forgetting.

Inspired by these works, we propose Look Forward to Walk Backward (LF2WB), a memory-enhanced perceptive locomotion controller that encodes forward egocentric depth

and proprioception into a compact, associative terrain memory using a delta-rule selective-update backbone DeltaNet. The policy conditions backward actions on retrieved memory and proprioception to anticipate obstacles without contact. During training we carry only a constant-size recurrent state across rollout windows and use chunkwise parallel updates inside each window. At deployment, LF2WB preserves terrain memory over extended horizons and executes agile backward traversals over steps and gaps as shown in Fig. 1. When the robot needs to retreat after forward traversal, it does not turn in place. It simply reverses the previous action sequence, and as long as the terrain was traversable in the forward direction, our memory framework provides accurate backward terrain information to guide the return.

Our contributions are as follows:

- We present Look Forward to Walk Backward (LF2WB), a single-stage training framework that learns to write terrain memory during forward motion and read it for collision-free backward locomotion without rearward vision.
- We integrate a delta-rule update mechanism that retains implicit terrain information, preventing uncontrolled memory accumulation while avoiding aggressive memory resets. This structure uses a constant-size recurrent state, supports chunkwise parallel training, and offers constant-time, constant-memory per-step inference with no KV cache, which makes it practical for low-cost onboard processors.
- We demonstrate zero-shot real-world deployment on a low-cost quadruped robot, achieving highly dynamic agile backward maneuvers over challenging step and gap terrains using only forward egocentric vision.

II. RELATED WORK

A. Learning-based Legged Locomotion

Recent learning-based blind controllers leverage reliable proprioception and achieve agile, robust motion in simulation and on hardware. They learn quickly, support online fine tuning, and withstand strong disturbances [26], [27], [28], [29], [30], [31], [32]. Other works enhance blind control with adaptation from privileged signals, or by training the policy together with a state estimator or implicit imagination modules [33], [34], [35], [36], [37], [38], [39]. Map-centric perceptive pipelines use elevation, traversability, or voxel representations to perceive terrain and have shown promising results. They often require well-tuned and heavy upstream modules to avoid latency, drift, and integration noise on fast or highly dynamic motions [40], [41], [42], [43], [44], [45], [6], [46]. Direct cross-modal fusion of forward vision with proprioception enables vigorous and agile motion on complex terrain [47], [48], [49], [50]. Adding memory modules, such as transformers or LSTM/RNNs, improves the use of historical information. This helps state estimation in partially observable settings and supports remembering terrain that is not currently visible underfoot [1], [2], [51], [13], [14], [15], [7], [5], [4], [52], [18], [17], [8].

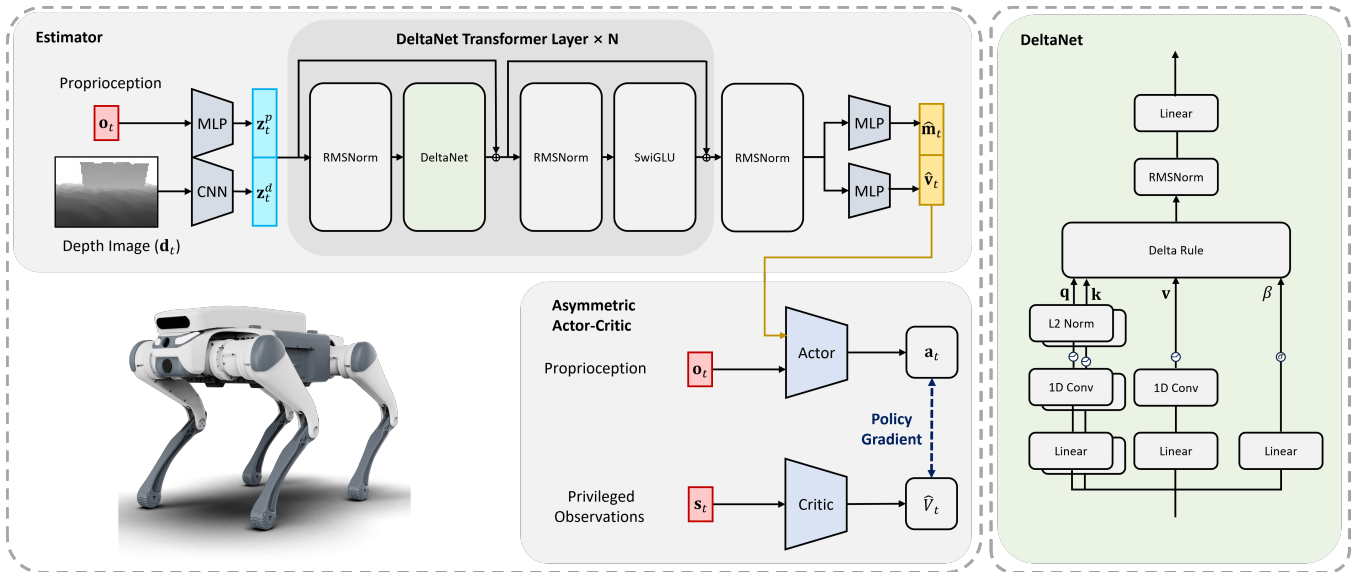


Fig. 2: Overview of the proposed LF2WB framework. We concurrently train a delta-rule estimator and an asymmetric actor-critic. The estimator learns a compact terrain memory with strong retrieval ability for backward motion, enabling the policy to execute dynamic backward maneuvers using only forward egocentric vision.

B. Linear-time Recurrent Sequence Models

Early linear-time formulations showed that causal attention can be rewritten as kernelized associative scans with recurrent state updates, reducing complexity to $O(L)$ for a sequence of length L while preserving a parallel training view [19]. This perspective links naturally to fast-weight memory: linearized attention programs a finite associative memory via outer-product updates, and delta-rule corrections address capacity and editability issues in purely additive schemes [24]. Building on this fast-weight view, delta-rule linear RNNs developed a remove-then-write update together with a Householder/WY-based sequence-parallel training algorithm that scales to long sequences while maintaining linear-time inference [25]. In parallel, retentive networks instantiated a retention operator with a constant-size recurrent state at inference and explicitly described parallel, recurrent, and chunkwise-recurrent computation forms [20]. Hybrid recurrent-Transformer designs pursued linear-time decoding at scale [53]. Structured state-space modeling introduced selective, input-dependent dynamics with hardware-aware selective-scan kernels for linear-time inference [22], and a subsequent unifying framework connected attention and SSMS via semiseparable operators with improved algorithms [54]. Within the linear-attention family, hardware-efficient gated variants provided input-conditioned gating and I/O-aware kernels under chunkwise-parallel schedules [21], and gated-delta formulations extended delta-rule models with data-dependent decay and parallel training refinements [55].

III. METHOD

A. Overview

We maintain an implicit internal memory constructed from forward egocentric sensing and proprioception, and use it throughout motion to support backward locomotion when rearward exteroception is unavailable. As shown in Fig. 2,

the system contains three modules trained concurrently on the same on-policy rollouts:

- 1) *Estimator*: It fuses a forward depth image $\mathbf{d}_t \in \mathbb{R}^{H_d \times W_d}$ and proprioception $\mathbf{o}_t \in \mathbb{R}^{D_o}$ to predict a body-centric elevation map $\hat{\mathbf{m}}_t \in \mathbb{R}^{H_m \times W_m}$ covering 1.1 m \times 2.5 m, and a base linear velocity estimate $\hat{\mathbf{v}}_t \in \mathbb{R}^{D_v}$.
- 2) *Actor*: π_θ : takes deployable observations $[\hat{\mathbf{m}}_t, \hat{\mathbf{v}}_t, \mathbf{o}_t]$ and outputs an action \mathbf{a}_t .
- 3) *Critic*: V_ϕ : takes privileged inputs $[\mathbf{m}_t, \mathbf{v}_t, \mathbf{o}_t]$ and predicts value V_t .

B. Estimator

1) *Encoders and Cross-Modal Fusion*: Depth and proprioception are encoded separately, then fused:

$$\mathbf{z}_t^{(d)} = f_{\text{cnn}}(\mathbf{d}_t) \in \mathbb{R}^{D_d}, \quad (1)$$

$$\mathbf{z}_t^{(p)} = f_{\text{mlp}}(\mathbf{o}_t) \in \mathbb{R}^{D_p}, \quad (2)$$

$$\mathbf{x}_t = [\mathbf{z}_t^{(d)}; \mathbf{z}_t^{(p)}] \in \mathbb{R}^{D_x}. \quad (3)$$

The fused token \mathbf{x}_t feeds a stack of L DeltaNet-Transformer layers.

2) *DeltaNet-Transformer Layer*: We follow a LLaMA-style block and replace self-attention with a DeltaNet layer [25]. Each layer uses RMSNorm pre-norm, applies normalization before the output projection, and is followed by a SwiGLU feed-forward network.

Let the layer input be $\mathbf{x}_t^{(\ell-1)} \in \mathbb{R}^d$ and $\tilde{\mathbf{x}}_t = \text{RMSNorm}(\mathbf{x}_t^{(\ell-1)})$. Feature maps (SiLU + L_2 normalization) define keys/queries as

$$\mathbf{k}_t^{(\ell)} = \frac{\text{SiLU}(\mathbf{W}_K^{(\ell)} \tilde{\mathbf{x}}_t)}{\|\text{SiLU}(\mathbf{W}_K^{(\ell)} \tilde{\mathbf{x}}_t)\|_2}, \quad (4)$$

$$\mathbf{q}_t^{(\ell)} = \frac{\text{SiLU}(\mathbf{W}_Q^{(\ell)} \tilde{\mathbf{x}}_t)}{\|\text{SiLU}(\mathbf{W}_Q^{(\ell)} \tilde{\mathbf{x}}_t)\|_2}. \quad (5)$$

The attention value vector (not the robot velocity) is

$$\mathbf{v}_{t,\text{att}}^{(\ell)} = \mathbf{W}_V^{(\ell)} \tilde{\mathbf{x}}_t \in \mathbb{R}^{d_v}, \quad (6)$$

and the overwrite scalar is

$$\beta_t^{(\ell)} = \sigma\left(\mathbf{w}_\beta^{(\ell)\top} \tilde{\mathbf{x}}_t\right) \in (0, 1). \quad (7)$$

The per-layer state matrix $\mathbf{S}_t^{(\ell)} \in \mathbb{R}^{d_v \times d_k}$ is updated by the Delta Rule (remove-then-write along $\mathbf{k}_t^{(\ell)}$):

$$\mathbf{S}_t^{(\ell)} = \mathbf{S}_{t-1}^{(\ell)} \left(\mathbf{I} - \beta_t^{(\ell)} \mathbf{k}_t^{(\ell)} \mathbf{k}_t^{(\ell)\top} \right) + \beta_t^{(\ell)} \mathbf{v}_{t,\text{att}}^{(\ell)} \mathbf{k}_t^{(\ell)\top}, \quad (8)$$

$$\mathbf{o}_{t,\text{att}}^{(\ell)} = \mathbf{S}_t^{(\ell)} \mathbf{q}_t^{(\ell)}. \quad (9)$$

With L_2 -normalized \mathbf{k}, \mathbf{q} , the eigenvalues of $\mathbf{I} - \beta \mathbf{k} \mathbf{k}^\top$ are $\{1\}^{d-1} \cup \{1-\beta\} \subset [0, 1]$, stabilizing training and interpreting $\beta=1$ as a projection that erases one subspace while retaining the other $d-1$.

Before output projection we apply RMSNorm:

$$\hat{\mathbf{o}}_{t,\text{att}}^{(\ell)} = \text{RMSNorm}(\mathbf{o}_{t,\text{att}}^{(\ell)}), \quad (10)$$

$$\Delta \mathbf{x}_t^{(\ell)} = \mathbf{W}_O^{(\ell)} \hat{\mathbf{o}}_{t,\text{att}}^{(\ell)}, \quad (11)$$

$$\mathbf{x}_t'^{(\ell)} = \mathbf{x}_t^{(\ell-1)} + \Delta \mathbf{x}_t^{(\ell)}. \quad (12)$$

SwiGLU FFN with pre-norm:

$$\tilde{\mathbf{x}}_t^{(\ell)} = \text{RMSNorm}(\mathbf{x}_t'^{(\ell)}), \quad (13)$$

$$\text{FFN}(\tilde{\mathbf{x}}_t^{(\ell)}) = \mathbf{W}_2^{(\ell)} \text{SwiGLU}(\mathbf{W}_1^{(\ell)} \tilde{\mathbf{x}}_t^{(\ell)}), \quad (14)$$

$$\mathbf{x}_t^{(\ell)} = \mathbf{x}_t'^{(\ell)} + \text{FFN}(\tilde{\mathbf{x}}_t^{(\ell)}). \quad (15)$$

The layer above replaces self-attention one-for-one; depth/width/FFN dimensions and residual layout match a standard Transformer block.

3) *Chunkwise Parallel Training*: For long sequences, we adopt a chunkwise parallel algorithm that vectorizes intra-chunk computation while carrying a fixed-size inter-chunk state. Split a length- L sequence into L/C chunks of size C . For chunk index u , stack the per-token projections and overwrite gates row-wise

$$\mathbf{Q}[u], \mathbf{K}[u], \mathbf{V}[u] \in \mathbb{R}^{C \times d_k}, \mathbb{R}^{C \times d_k}, \mathbb{R}^{C \times d_v}, \quad (16)$$

$$\beta[u] \in \mathbb{R}^C. \quad (17)$$

Using the UT transform to expose large matmuls, define

$$\mathbf{T}[u] = \left(\mathbf{I} + \text{strictLower}(\text{diag}(\beta[u]) \mathbf{K}[u] \mathbf{K}[u]^\top) \right)^{-1} \text{diag}(\beta[u]), \quad (18)$$

$$\mathbf{W}[u] = \mathbf{T}[u] \mathbf{K}[u], \quad (19)$$

$$\mathbf{U}[u] = \mathbf{T}[u] \mathbf{V}[u]. \quad (20)$$

Let $\mathbf{S}[u] \in \mathbb{R}^{d_v \times d_k}$ be the (small) state entering chunk u . The cross-chunk update and the chunk outputs are

$$\mathbf{S}[u+1] = \mathbf{S}[u] \left(\mathbf{I} - \mathbf{K}[u]^\top \mathbf{W}[u] \right) + \mathbf{U}[u]^\top \mathbf{K}[u], \quad (21)$$

$$\begin{aligned} \mathbf{O}[u] &= \left(\mathbf{Q}[u] \mathbf{K}[u]^\top \odot \mathbf{M}_C \right) \left(\mathbf{U}[u] - \mathbf{W}[u] \mathbf{S}[u]^\top \right) \\ &\quad + \mathbf{Q}[u] \mathbf{S}[u]^\top, \end{aligned} \quad (22)$$

where \mathbf{M}_C is the strict lower-triangular causal mask. All intra-chunk interactions are parallel and GPU-friendly; only the compact $\mathbf{S}[u]$ is propagated across chunks. No per-token hidden histories are materialized. The training complexity is $O(LCd + Ld^2)$, while test-time inference is strictly recurrent with *constant* per-head state $\mathbf{S}_t^{(\ell)}$.

4) *Heads and Supervised Losses*: The top layer feature $\mathbf{x}_t^{(L)}$ is decoded to

$$\hat{\mathbf{m}}_t \in \mathbb{R}^{H_m \times W_m}, \quad (23)$$

$$\hat{\mathbf{v}}_t \in \mathbb{R}^{D_v}, \quad (24)$$

We minimize

$$\mathcal{L}_{\text{est}} = \lambda_m \|\hat{\mathbf{m}}_t - \mathbf{m}_t\|_2^2 + \lambda_v \|\hat{\mathbf{v}}_t - \mathbf{v}_t\|_2^2. \quad (25)$$

where $\lambda_m, \lambda_v > 0$ are scalar weights, \mathbf{m}_t is the ground-truth elevation map, and \mathbf{v}_t is the ground-truth base linear velocity.

C. Asymmetric Actor-Critic

We adopt the asymmetric actor-critic framework [56]. Only deployable signals serve as the input of the actor and privileged inputs are used to stabilize value learning:

$$\tilde{\mathbf{s}}_t = [\text{vec}(\hat{\mathbf{m}}_t); \hat{\mathbf{v}}_t; \mathbf{o}_t], \quad (26)$$

$$\mathbf{a}_t \sim \pi_\theta(\cdot | \tilde{\mathbf{s}}_t), \quad (27)$$

$$\mathbf{s}_t = [\text{vec}(\mathbf{m}_t); \mathbf{v}_t; \mathbf{o}_t], \quad (28)$$

$$V_t = V_\phi(\mathbf{s}_t). \quad (29)$$

The actor is deployable; the critic uses privileged inputs for training only.

We use PPO for on-policy reinforcement learning [57]. Let $r_t(\theta) = \frac{\pi_\theta(\mathbf{a}_t | \tilde{\mathbf{s}}_t)}{\pi_{\theta_{\text{old}}}(\mathbf{a}_t | \tilde{\mathbf{s}}_t)}$ be the importance ratio, \hat{A}_t the GAE advantage, and R_t the return target.

$$r_t(\theta) = \frac{\pi_\theta(\mathbf{a}_t | \tilde{\mathbf{s}}_t)}{\pi_{\theta_{\text{old}}}(\mathbf{a}_t | \tilde{\mathbf{s}}_t)}, \quad (30)$$

$$\mathcal{L}_{\text{clip}}(\theta) = \mathbb{E} \left[\min \left(r_t \hat{A}_t, \text{clip}(r_t, 1-\epsilon, 1+\epsilon) \hat{A}_t \right) \right], \quad (31)$$

$$\mathcal{L}_V(\phi) = \mathbb{E} \left[(V_\phi(\mathbf{s}_t) - R_t)^2 \right], \quad (32)$$

$$\mathcal{L}_{\text{ent}}(\theta) = \mathbb{E} \left[\mathcal{H}(\pi_\theta(\cdot | \tilde{\mathbf{s}}_t)) \right], \quad (33)$$

where $\mathcal{H}(\cdot)$ is entropy and $\text{clip}(\cdot)$ is applied elementwise. The minimized PPO objective is

$$\mathcal{L}_{\text{PPO}}(\theta, \phi) = -\mathcal{L}_{\text{clip}}(\theta) + c_V \mathcal{L}_V(\phi) - c_H \mathcal{L}_{\text{ent}}(\theta), \quad (34)$$

with scalar weights $c_V, c_H > 0$.

D. Training Details

We optimize the estimator, actor, and critic concurrently, which avoids pretraining and exposes the policy to realistic estimation noise [35], [36]:

$$\mathcal{L} = \mathcal{L}_{\text{PPO}} + \alpha \mathcal{L}_{\text{est}}. \quad (35)$$

with $\alpha > 0$.

In the massively parallel on-policy setting, we collect short contiguous rollout windows of K steps per environment so that the update batch $B = N_{\text{env}} \times K$ remains in a regime

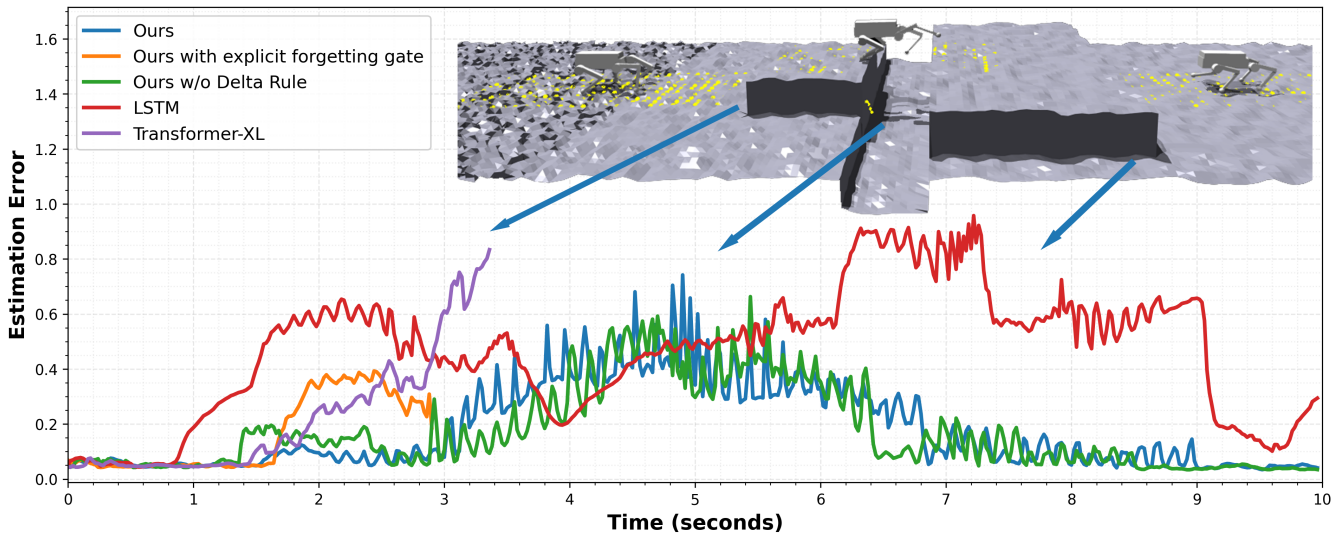


Fig. 3: RMSE of the body-centric elevation estimate along the executed path. The robot starts on the left of the image and walks backward to the right; 0 s corresponds to the leftmost position where the backward phase begins. Curve locations at terrain discontinuities are marked by arrows. Once the robot falls, we stop recording error.

that preserves temporal coherence while sustaining GPU throughput [29]. Episodes span multiple rollout updates, and we carry only the estimator’s compact recurrent state $\mathbf{S}_t^{(\ell)}$ across windows.

The first state of the next rollout window is detached, reducing memory. Crucially, compared to self-attention Transformers, the size of each $\mathbf{S}_t^{(\ell)}$ is constant and does not grow with the desired memory horizon.

During rollout interaction with the simulator to collect data, the estimator runs in serial recurrent mode, updating the state $\mathbf{S}_t^{(\ell)}$ step by step. After a window of K steps is collected, the optimization phase replays the window and computes gradients with the chunkwise parallel algorithm for the estimator.

In our setup, the maximum episode length is set to 20 s. Accordingly we aim for a 20 s terrain memory. We set $K=320$ with a 50 Hz control rate, so each rollout window covers 6.4 s. During training we command the robot to move forward for 10 s and then backward for 10 s, over a curriculum that progressively increases the difficulty of high platforms and gaps, as well as their random combinations. Other training configurations such as reward design, domain randomization, and PPO hyperparameters follow [7] exactly.

IV. EXPERIMENTS

A. Comparative Experiments

To systematically investigate performance across complex terrains, we evaluate in Isaac Gym using the DEEP Robotics Lite3 quadruped robot and compare estimator backbones under identical overall structure, PPO settings, reward, and domain randomization; only the estimator differs. The four methods that we compared ours against are:

- **Ours with explicit forgetting gate (Gated DeltaNet [55]):** Augments the Delta Rule update with a learned input-conditioned decay gate that

down-weights the retained state before writing new content for rapid purging of spurious associations.

- **Ours w/o Delta Rule (Linear Attention [19]):** Removes the Delta Rule and keeps only additive key-value accumulation. It preserves constant-size recurrent state at inference but lacks a principled erasure mechanism.
- **LSTM [58]:** A stacked LSTM with matched network size, trained with the same on-policy pipeline and TBPTT. It offers fixed-size recurrent state and serial inference.
- **Transformer-XL [16]:** A decoder-only Transformer that reuses cached detached hidden states from preceding segments when training.

Because Transformer-XL must retain preceding segment memories via cached hidden states whose size grows with the desired memory length and thus increases GPU memory usage, its training memory window is limited to 240 steps (4.8 s) versus 320 steps (6.4 s) for all other backbones.

To examine backward elevation estimation error, each policy controls the robot to walk forward 10 s and then backward 10 s. As shown in Fig. 3, during the backward phase we plot the RMSE of the body-centric elevation estimate along the executed path. When the estimator’s error becomes large enough that the policy can no longer negotiate the upcoming terrain and fails, the run terminates and its error trace naturally truncates at that time. From these traces, our method exhibits markedly lower error on complex terrain, supporting collision-free backward maneuvers. **Ours w/o Delta Rule** also performs well on this in-distribution 20 s task, with accuracy comparable to **Ours**. **LSTM** produces less accurate estimates, yet concurrent training enables the actor to tolerate moderate bias/noise and still complete the trail. **Ours w/ explicit forgetting gate** and **Transformer-XL** show large bias and typically fail. Although self-attention is powerful, **Transformer-XL**’s effective memory is curtailed by its short segment window under our budget, preventing

TABLE I: Success rate (%) across difficulty $\delta \in \{0.2, 0.4, 0.6, 0.8\}$ and protocol **P1–P5**.

Estimator Backbone	P1 (F10s→B10s)				P2 (F20s→B10s)				P3 (F20s→B20s)				P4 (F30s→B10s)				P5 (F30s→B30s)			
	0.2	0.4	0.6	0.8	0.2	0.4	0.6	0.8	0.2	0.4	0.6	0.8	0.2	0.4	0.6	0.8	0.2	0.4	0.6	0.8
Ours	90.1	86.9	81.7	75.0	89.5	83.5	77.7	71.7	49.0	42.2	37.9	30.4	87.0	83.6	78.9	72.1	16.1	10.9	10.8	10.5
Ours with explicit forgetting gate	70.1	66.4	57.5	46.7	72.7	67.4	59.0	46.6	25.3	19.1	10.5	8.8	66.1	60.8	56.5	46.5	9.8	5.1	3.3	3.4
Ours w/o Delta Rule	85.1	86.8	81.5	68.9	62.0	54.9	42.0	30.7	36.7	32.8	23.9	12.4	39.6	35.1	22.8	20.0	7.0	6.2	3.0	2.7
LSTM	57.1	45.0	32.8	25.7	60.3	48.0	31.3	25.7	20.6	7.6	4.3	3.8	64.3	52.2	35.8	29.7	5.8	2.3	1.0	1.1
Transformer-XL	28.9	26.6	22.7	16.7	30.4	24.7	20.1	16.2	7.8	7.7	3.9	1.9	28.2	24.7	19.7	15.5	2.1	1.2	1.1	1.0

learning of longer-horizon terrain memory.

We then evaluate all backbones across four terrain difficulty levels $\delta \in \{0.2, 0.4, 0.6, 0.8\}$ with increasing complexity. Difficulty is defined by linearly scaling the maxima used during training: the hardest training terrain contains random mixtures of steps up to 0.8 m in height and gaps up to 1.1 m in width, combined with additional vertical offsets (e.g., a 1.0 m gap whose two sides differ in height by 0.7 m). A test trail at difficulty δ uses the same random mixing but scaled linearly, i.e., $\text{step}(\delta) = 0.8\delta$ m and $\text{gap}(\delta) = 1.1\delta$ m; for example, $\delta=0.8$ yields 0.64 m steps and 0.88 m gaps.

We consider five forward/backward horizon protocols at each difficulty:

- 1) **P1 (F10s→B10s)**: In-distribution evaluation that matches the training schedule. This setting tests baseline backward recall.
- 2) **P2 (F20s→B10s)**: Capacity/overload check. An extra 10s of forward writing examines whether irrelevant or erroneous content can be removed without harming backward performance.
- 3) **P3 (F20s→B20s)**: Long-horizon recall. This setting assesses whether distant yet relevant terrain information remain retrievable when the backward horizon is doubled relative to training.
- 4) **P4 (F30s→B10s)**: Strong overload with a short backward retrieval. Write capacity is pushed substantially beyond training while the backward horizon remains short.
- 5) **P5 (F30s→B20s)**: Strong overload with a long backward retrieval. This setting stresses maximum write load together with extended backward retrieval.

As shown in Table I, **Ours** achieves the best performance across all protocol groups and difficulty levels. Although performance decreases as the mismatch between the test horizon and the training horizon grows, our method still exhibits markedly stronger generalization than other baselines. We attribute this to the delta-rule selective update, which improves in-context recall. Training with longer memory horizons is expected to further mitigate the performance drop. Due to the explicit forgetting gate and the update scheme during training that does not optimize over the entire sequence (similar to TBPTT), **Ours with explicit forgetting gate** and **LSTM** have difficulty maintaining long-range information, leading to large drops on **P3** and **P5**. On the other hand,

the forgetting gate can quickly remove redundant content in the memory state, so the extra forward time in **P2** and **P4** has limited impact early in the backward segment. **Ours w/o Delta Rule** performs well on the training-like setting but has no removal mechanism, continuously accumulating key-value outer products in a fixed-size state whose finite capacity eventually leads to overload, producing a sharper decline on **P2** and **P4**. Conversely, because it does not forget, it can preserve older information better than gated models on **P3**, where long-range retrieval dominates. **Transformer-XL** underperforms due to the limited memory segment window enforced during training.

B. Real-World Deployment

We deploy the controller on a DEEP Robotics Lite3 low-cost quadruped equipped with an onboard RK3588 processor and a forward-facing RealSense depth camera. Images are streamed at 10 Hz and the control loop runs at 50 Hz. The policy executes strictly recurrent inference with a fixed-size memory state and constant per-step computation, which keeps the hardware control loop stable.

As shown in Fig. 4, we evaluate on multiple terrains, including a 0.6 m high step, a 0.6 m wide gap, and composite layouts that combine steps and gaps. The longest trails are close to 8 m, which is nearly the limit of our test area. In Fig. 4(a), the robot jumps backward across a gap with a 0.3 m height drop between two edges, which is rarely demonstrated even in forward perceptive locomotion by prior works. Our system not only performs precise forward jumps with perception, but also executes the same maneuver in reverse by relying on accurate memory. In Fig. 4(b), the robot jumps backward onto a high step. In Fig. 4(c), we show the full pipeline of first moving forward to memorize the surroundings and then moving backward, climbing multiple obstacles while retrieving the correct information from the memory state. Despite cumulative real-world noise, the memory module remains robust and adaptive.

V. CONCLUSION

In this work, we present Look Forward to Walk Backward (LF2WB), a single-stage framework that uses only forward vision to build and maintain an efficient terrain memory and retrieve it for collision-free backward locomotion without rearward sensing. A delta-rule selective update keeps this memory compact, while chunkwise-parallel training makes


(a) Move Backward



(b) Move Backward



(c) Move Forward, Then Backward



Moving Direction ➔

Fig. 4: Real-world deployment on a robot equipped with a forward-facing depth camera. For clarity, all sequences are visualized such that the robot moves from left to right.

learning efficient. Serial recurrent inference with constant-size states makes the method well-suited for efficient deployment on onboard processors. In our experiments, LF2WB outperforms other recurrent networks and Transformer-XL under varied forward/backward horizon stresses. We further demonstrate zero-shot deployment on a DEEP Robotics Lite3 quadruped robot, which relies solely on a forward depth camera and the learned memory to execute agile backward maneuvers over challenging composite terrains that combine high steps and wide gaps.

However, certain limitations still exist. The memory is implicit, so its effective horizon is not formally guaranteed. We sometimes observe surprisingly strong generalization. For example, although we train only forward and backward motion, the robot still shows some memory capability when turning. But we also see occasional failures on seemingly simple cases. We believe broader task diversity and scene coverage during training will improve reliability. Incorporating estimator uncertainty into the concurrent training loop could further guide the actor to respect the estimator’s capability limits and enhance overall robustness.

REFERENCES

- [1] Z. Zhuang, Z. Fu, J. Wang, C. Atkeson, S. Schwertfeger, C. Finn, and H. Zhao, “Robot parkour learning,” *arXiv preprint arXiv:2309.05665*, 2023.
- [2] X. Cheng, K. Shi, A. Agarwal, and D. Pathak, “Extreme parkour with legged robots,” *arXiv preprint arXiv:2309.14341*, 2023.
- [3] E. Chane-Sane, J. Amigo, T. Flayols, L. Righetti, and N. Mansard, “Soloparkour: Constrained reinforcement learning for visual locomotion from privileged experience,” *arXiv preprint arXiv:2409.13678*, 2024.
- [4] H. Lai, J. Cao, J. Xu, H. Wu, Y. Lin, T. Kong, Y. Yu, and W. Zhang, “World model-based perception for visual legged locomotion,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 11 531–11 537.
- [5] S. Li, S. Luo, J. Wu, and Q. Zhu, “Move: Multi-skill omnidirectional legged locomotion with limited view in 3d environments,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 7647–7653.
- [6] Z. Wang, Y. Li, L. Xu, H. Shi, Z. Ma, Z. Chu, C. Li, F. Gao, K. Yang, and K. Wang, “Sf-tim: A simple framework for enhancing quadrupedal robot jumping agility by combining terrain imagination and measurement,” *arXiv preprint arXiv:2408.00486*, 2024.
- [7] S. Luo, S. Li, R. Yu, Z. Wang, J. Wu, and Q. Zhu, “Pie: Parkour with implicit-explicit learning framework for legged robots,” *IEEE Robotics and Automation Letters*, 2024.
- [8] Y. Zhang, Q. Qian, T. Hou, P. Zhai, X. Wei, K. Hu, J. Yi, and L. Zhang, “Renet: Fault-tolerant motion control for quadruped robots via redundant estimator networks under visual collapse,” *IEEE Robotics and Automation Letters*, pp. 1–8, 2025.
- [9] R. J. Williams and J. Peng, “An efficient gradient-based algorithm for on-line training of recurrent network trajectories,” *Neural computation*, vol. 2, no. 4, pp. 490–501, 1990.
- [10] R. Pascanu, T. Mikolov, and Y. Bengio, “On the difficulty of training recurrent neural networks,” in *International conference on machine learning*. Pmlr, 2013, pp. 1310–1318.
- [11] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [12] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei, “Scaling laws for neural language models,” *arXiv preprint arXiv:2001.08361*, 2020.

- [13] I. Radosavovic, T. Xiao, B. Zhang, T. Darrell, J. Malik, and K. Sreenath, "Learning humanoid locomotion with transformers," *CoRR*, 2023.
- [14] —, "Real-world humanoid locomotion with reinforcement learning," *Science Robotics*, vol. 9, no. 89, p. eadi9579, 2024.
- [15] I. Radosavovic, S. Kamat, T. Darrell, and J. Malik, "Learning humanoid locomotion over challenging terrain," *arXiv preprint arXiv:2410.03654*, 2024.
- [16] Z. Dai, Z. Yang, Y. Yang, J. Carbonell, Q. V. Le, and R. Salakhutdinov, "Transformer-xl: Attentive language models beyond a fixed-length context," *arXiv preprint arXiv:1901.02860*, 2019.
- [17] H. Zheng, Y. Cheng, H. Liu, H. Liu, L. Ye, X. Wang, and B. Liang, "Huwo: Building physical interaction world models for humanoid robot locomotion."
- [18] K.-Y. Lin and S. X. Yu, "Let humanoids hike! integrative skill development on complex trails," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 22498–22507.
- [19] A. Katharopoulos, A. Vyas, N. Pappas, and F. Fleuret, "Transformers are rns: Fast autoregressive transformers with linear attention," in *International conference on machine learning*. PMLR, 2020, pp. 5156–5165.
- [20] Y. Sun, L. Dong, S. Huang, S. Ma, Y. Xia, J. Xue, J. Wang, and F. Wei, "Retentive network: A successor to transformer for large language models," *arXiv preprint arXiv:2307.08621*, 2023.
- [21] S. Yang, B. Wang, Y. Shen, R. Panda, and Y. Kim, "Gated linear attention transformers with hardware-efficient training," *arXiv preprint arXiv:2312.06635*, 2023.
- [22] A. Gu and T. Dao, "Mamba: Linear-time sequence modeling with selective state spaces," *arXiv preprint arXiv:2312.00752*, 2023.
- [23] B. Widrow and M. E. Hoff, "Adaptive switching circuits," in *Neurocomputing: foundations of research*, 1988, pp. 123–134.
- [24] I. Schlag, K. Irie, and J. Schmidhuber, "Linear transformers are secretly fast weight programmers," in *International conference on machine learning*. PMLR, 2021, pp. 9355–9366.
- [25] S. Yang, B. Wang, Y. Zhang, Y. Shen, and Y. Kim, "Parallelizing linear transformers with the delta rule over sequence length," *Advances in neural information processing systems*, vol. 37, pp. 115491–115522, 2024.
- [26] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter, "Learning agile and dynamic motor skills for legged robots," *Science Robotics*, vol. 4, no. 26, p. eaau5872, 2019.
- [27] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning quadrupedal locomotion over challenging terrain," *Science robotics*, vol. 5, no. 47, p. eabc5986, 2020.
- [28] Z. Fu, A. Kumar, J. Malik, and D. Pathak, "Minimizing energy consumption leads to the emergence of gaits in legged robots," *arXiv preprint arXiv:2111.01674*, 2021.
- [29] 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.
- [30] L. Smith, J. C. Kew, X. B. Peng, S. Ha, J. Tan, and S. Levine, "Legged robots that keep on learning: Fine-tuning locomotion policies in the real world," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 1593–1599.
- [31] G. B. Margolis, G. Yang, K. Paigwar, T. Chen, and P. Agrawal, "Rapid locomotion via reinforcement learning," *The International Journal of Robotics Research*, p. 02783649231224053, 2022.
- [32] 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.
- [33] A. Kumar, Z. Fu, D. Pathak, and J. Malik, "Rma: Rapid motor adaptation for legged robots," *arXiv preprint arXiv:2107.04034*, 2021.
- [34] W. Yu, J. Tan, Y. Bai, E. Coumans, and S. Ha, "Learning fast adaptation with meta strategy optimization," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2950–2957, 2020.
- [35] 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.
- [36] 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.
- [37] J. Long, Z. Wang, Q. Li, J. Gao, L. Cao, and J. Pang, "Hybrid internal model: A simple and efficient learner for agile legged locomotion," *arXiv preprint arXiv:2312.11460*, 2023.
- [38] H. Wang, H. Luo, W. Zhang, and H. Chen, "Cts: Concurrent teacher-student reinforcement learning for legged locomotion," *IEEE Robotics and Automation Letters*, 2024.
- [39] Z. Luo, Y. Dong, X. Li, R. Huang, Z. Shu, E. Xiao, and P. Lu, "Moral: Learning morphologically adaptive locomotion controller for quadrupedal robots on challenging terrains," *IEEE Robotics and Automation Letters*, 2024.
- [40] P. Fankhauser, M. Bloesch, and M. Hutter, "Probabilistic terrain mapping for mobile robots with uncertain localization," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3019–3026, 2018.
- [41] Y. Pan, X. Xu, Y. Wang, X. Ding, and R. Xiong, "Gpu accelerated real-time traversability mapping," in *2019 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. IEEE, 2019, pp. 734–740.
- [42] T. Miki, L. Wellhausen, R. Grandia, F. Jenelten, T. Homberger, and M. Hutter, "Elevation mapping for locomotion and navigation using gpu," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 2273–2280.
- [43] D. Hoeller, N. Rudin, C. Choy, A. Anandkumar, and M. Hutter, "Neural scene representation for locomotion on structured terrain," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 8667–8674, 2022.
- [44] 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.
- [45] D. Hoeller, N. Rudin, D. Sako, and M. Hutter, "Anymal parkour: Learning agile navigation for quadrupedal robots," *Science Robotics*, vol. 9, no. 88, p. eadi7566, 2024.
- [46] J. He, C. Zhang, F. Jenelten, R. Grandia, M. Bächer, and M. Hutter, "Attention-based map encoding for learning generalized legged locomotion," *Science Robotics*, vol. 10, no. 105, p. eadv3604, 2025.
- [47] A. Loquercio, A. Kumar, and J. Malik, "Learning visual locomotion with cross-modal supervision," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 7295–7302.
- [48] Z. Fu, A. Kumar, A. Agarwal, H. Qi, J. Malik, and D. Pathak, "Coupling vision and proprioception for navigation of legged robots," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 17273–17283.
- [49] S. Kareer, N. Yokoyama, D. Batra, S. Ha, and J. Truong, "Vinl: Visual navigation and locomotion over obstacles," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 2018–2024.
- [50] D. Wang, X. Wang, X. Liu, J. Shi, Y. Zhao, C. Bai, and X. Li, "More: Mixture of residual experts for humanoid lifelike gaits learning on complex terrains," *arXiv preprint arXiv:2506.08840*, 2025.
- [51] A. Agarwal, A. Kumar, J. Malik, and D. Pathak, "Legged locomotion in challenging terrains using egocentric vision," in *Conference on robot learning*. PMLR, 2023, pp. 403–415.
- [52] Y. Yang, G. Shi, C. Lin, X. Meng, R. Scalise, M. G. Castro, W. Yu, T. Zhang, D. Zhao, J. Tan *et al.*, "Agile continuous jumping in discontinuous terrains," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 10245–10252.
- [53] B. Peng, E. Alcaide, Q. Anthony, A. Albalak, S. Arcadinho, S. Biderman, H. Cao, X. Cheng, M. Chung, M. Grella *et al.*, "Rwkv: Reinventing rns for the transformer era," *arXiv preprint arXiv:2305.13048*, 2023.
- [54] T. Dao and A. Gu, "Transformers are ssms: Generalized models and efficient algorithms through structured state space duality," *arXiv preprint arXiv:2405.21060*, 2024.
- [55] S. Yang, J. Kautz, and A. Hatamizadeh, "Gated delta networks: Improving mamba2 with delta rule," *arXiv preprint arXiv:2412.06464*, 2024.
- [56] L. Pinto, M. Andrychowicz, P. Welinder, W. Zaremba, and P. Abbeel, "Asymmetric actor critic for image-based robot learning," *arXiv preprint arXiv:1710.06542*, 2017.
- [57] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [58] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.