

GeoLoco: Leveraging 3D Geometric Priors from Visual Foundation Model for Robust RGB-Only Humanoid Locomotion

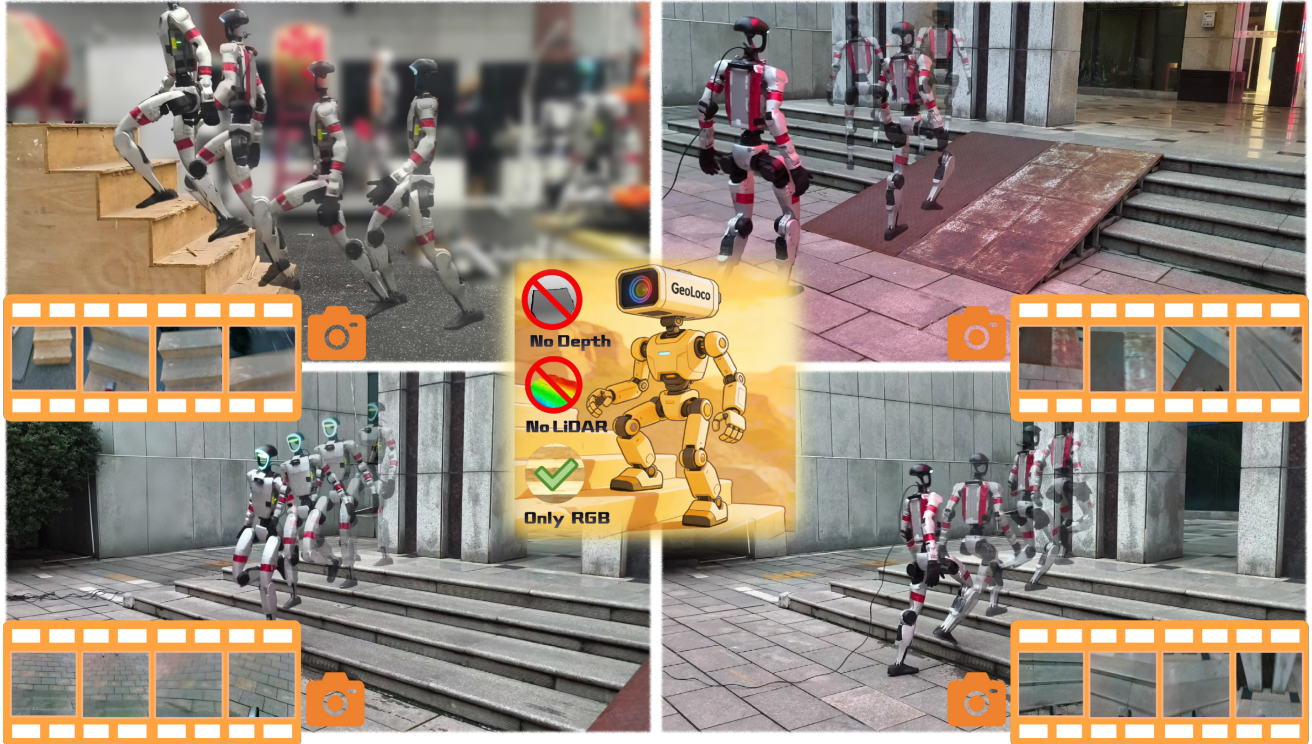


Fig. 1: **Overview of GeoLoco.** Our framework enables robust humanoid locomotion on diverse terrains—including stairs, ramps, and uneven blocks—using **only a monocular RGB camera**. By conceptualizing 2D pixels as high-dimensional 3D geometric representations, GeoLoco eliminates the dependency on active depth sensors (LiDAR/Depth) while maintaining expert-level traversal performance.

Abstract—The prevailing paradigm of perceptive humanoid locomotion relies heavily on active depth sensors. However, this depth-centric approach fundamentally discards the rich semantic and dense appearance cues of the visual world, severing low-level control from the high-level reasoning essential for general embodied intelligence. While monocular RGB offers a ubiquitous, information-dense alternative, end-to-end reinforcement learning from raw 2D pixels suffers from extreme sample inefficiency and catastrophic sim-to-real collapse due to the inherent loss of geometric scale. To break this deadlock, we propose GeoLoco, a purely RGB-driven locomotion framework that conceptualizes monocular images as high-dimensional 3D latent representations by harnessing the powerful geometric priors of a frozen, scale-aware Visual Foundation Model (VFM). Rather than naive feature concatenation, we design a proprioceptive-query multi-head cross-attention mechanism that dynamically attends to task-critical topological features conditioned on the robot’s real-time gait phase. Crucially, to prevent the policy from overfitting to superficial textures, we introduce a dual-head auxiliary learning scheme. This explicit regularization forces the high-dimensional latent space to strictly align with the physical terrain geometry, ensuring robust zero-shot sim-

to-real transfer. Trained exclusively in simulation, GeoLoco achieves robust zero-shot transfer to the Unitree G1 humanoid and successfully negotiates challenging terrains.

I. INTRODUCTION

Humanoid locomotion represents a fundamental pillar of general embodied intelligence, requiring the seamless integration of high-level perception and low-level motor control. Recently, deep reinforcement learning (DRL) has significantly accelerated this field, enabling robots to exhibit remarkable agility and gait emergence within high-fidelity simulations [1], [2]. However, bridging the gap to unstructured real-world environments remains a formidable challenge. While proprioception-only policies have demonstrated impressive disturbance rejection on flat or slightly uneven ground [3], [4], they remain fundamentally blind to upcoming geometric discontinuities such as stairs and discrete obstacles. Negotiating such terrains necessitates perceptive locomotion, where the control policy must actively reason

about terrain geometry to proactively adjust foot placement and whole-body posture.

The prevailing paradigm for perceptive locomotion heavily relies on active depth sensors, such as LiDAR or RGB-D cameras, to explicitly reconstruct local elevation maps or height fields [5], [6]. Although effective for geometric stability, this depth-centric paradigm is increasingly becoming an information silo that is misaligned with the broader trend of integrated embodied intelligence. By stripping away dense appearance and semantic cues in favor of sparse geometric representations, these methods disconnect the low-level locomotion policy from the rich contextual information required for cross-task reasoning and Vision-Language-Action (VLA) alignment. This decoupling hinders the potential for locomotion policies to benefit from the scaling laws of modern Visual Foundation Models (VFMs). Consequently, transitioning to a purely RGB-driven paradigm inherently preserves this dense contextual richness, establishing a scalable prerequisite for integrating low-level locomotion into general-purpose VLA frameworks.

Despite the potential of the RGB modality, achieving robust geometry-aware locomotion from a single monocular camera remains a formidable challenge. Standard 2D pixel observations inherently lack metric scale and are highly susceptible to visual ambiguities and lighting variations, making end-to-end RL notoriously sample-inefficient and brittle during simulation-to-real transfer (sim-to-real). However, the emergence of large-scale VFMs offers a transformative perspective. In this work, we posit that monocular RGB input should be re-conceptualized not merely as a 2D pixel array, but as a latent representation of the 3D world, unlocked by the scaling laws and geometric priors embedded in modern VFMs. By leveraging these pre-trained models, we can extract robust, metric-aware geometric cues that were fundamentally inaccessible to Convolutional Neural Network (CNN) baselines.

To address these challenges, we propose GeoLoco, a purely RGB-driven reinforcement learning framework for robust humanoid locomotion. Bypassing active depth sensors, GeoLoco employs a frozen, scale-aware visual geometry encoder to project monocular RGB observations into high-dimensional geometric features. As shown in Fig. 2, a lightweight cross-attention module fuses visual embeddings with proprioceptive observation. Guided by real-time kinematics, the policy dynamically focuses on critical terrain features like stair edges and riser heights. Crucially, to mitigate overfitting to superficial visual textures, we introduce a dual-head auxiliary learning scheme. This explicit regularization physically grounds the high-dimensional latent space by enforcing strict alignment with local terrain map and robot states. Trained exclusively in simulation, GeoLoco achieves zero-shot deployment on the Unitree G1 humanoid, successfully traversing severe geometric discontinuities without any real-world fine-tuning.

The main contributions of this paper are as follows:

- We propose a purely RGB-driven humanoid locomotion framework that bridges geometric precision with

semantic richness by conceptualizing monocular RGB as a 3D latent representation, effectively eliminating the reliance on active depth sensors.

- We introduce a lightweight visual-proprioceptive fusion architecture utilizing cross-attention, which efficiently integrates high-dimensional visual priors with proprioceptive states to facilitate geometry-aware whole-body control.
- We design a dual-head auxiliary representation learning scheme that regularizes the latent space by reconstructing terrain topography and predicting system dynamics, ensuring robust zero-shot sim-to-real transfer under diverse environmental conditions.

II. RELATED WORK

A. Bottlenecks in Perceptive Locomotion

While DRL has significantly enhanced humanoid agility [1], proprioceptive-only policies [3], [4] remain fundamentally blind to geometric discontinuities. Consequently, *perceptive locomotion* relying on active depth sensors has become indispensable [7], [5]. However, this depth-centric paradigm creates a critical information silo. By discarding dense appearance and semantic cues, it severs low-level motor control from the high-level reasoning essential for general embodied intelligence. Transitioning to a purely RGB-driven modality is therefore inevitable for seamless Vision-Language-Action (VLA) alignment.

B. Visual Representations: From 2D Priors to 3D Priors

Deriving metric geometry from monocular RGB remains a long-standing challenge [8]. Yet, extracting reliable, scale-aware geometry from monocular RGB remains a formidable bottleneck, explicitly necessitating a paradigm shift in visual representation. Recently, Visual Foundation Models (VFMs) have emerged as powerful feature extractors for robotics, driving advancements in humanoid whole-body control [9], general-purpose manipulation [10], and Vision-Language Navigation (VLN) [11].

However, these applications predominantly focus on high-level semantics, task instructions, or quasi-static behaviors. For high-frequency perceptive locomotion over complex terrains, the critical bottlenecks are real-time inference latency and the acquisition of metric-scale geometry.

C. Sim-to-Real Transfer for Visual Policy

The primary obstacle to RGB-only control is the substantial sim-to-real gap in visual appearance. Recently, pure pixel-to-action RL has emerged as a promising paradigm for humanoid control. Pioneering works such as VIRAL [12] and related policy transfers [13] have demonstrated impressive visual sim-to-real capabilities at scale. Furthermore, end-to-end visual RL has even been shown to facilitate emergent active perception in simulated humanoids [14].

However, prevailing RGB-driven methods typically train visual encoders entirely from scratch, necessitating aggressive Domain Randomization (DR) that exacerbates sample

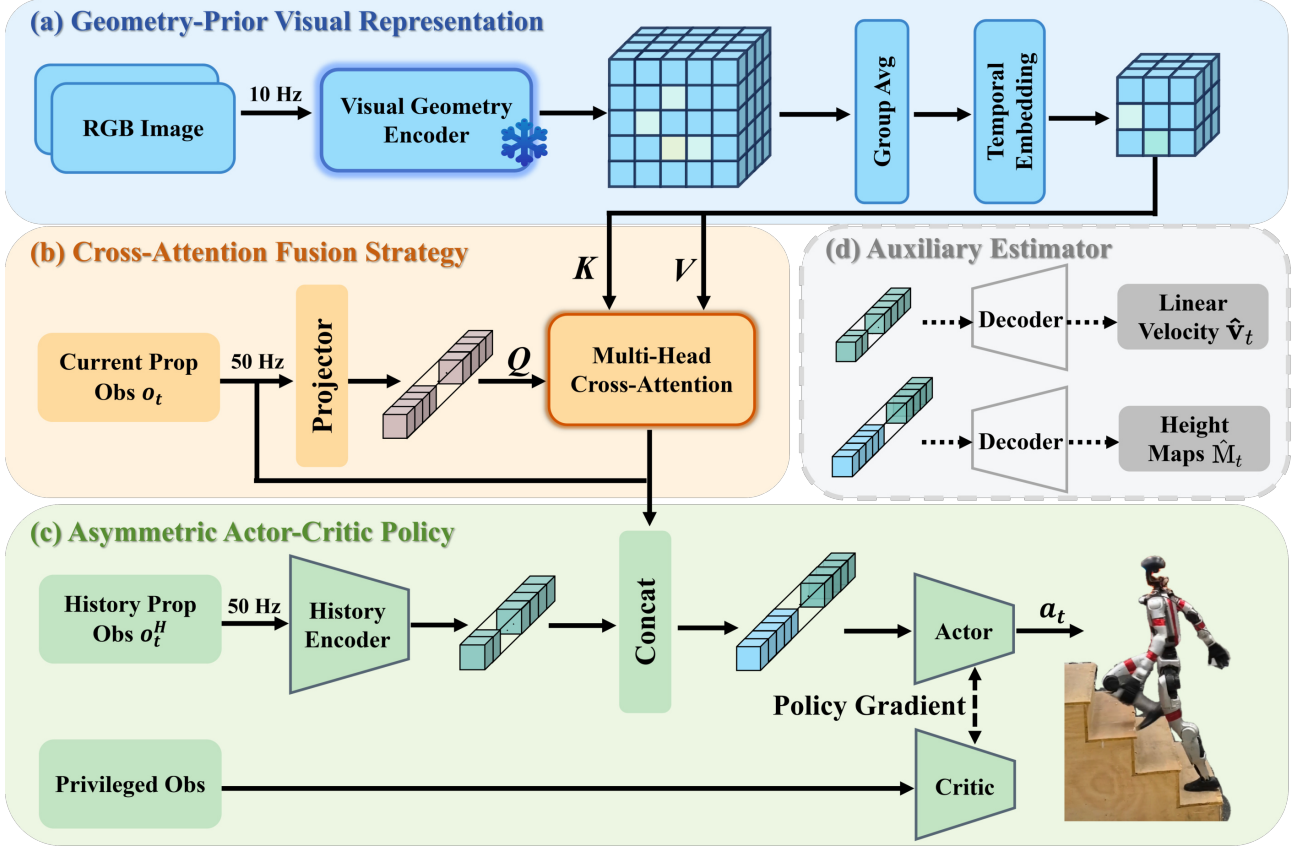


Fig. 2: **The GeoLoco Architecture.** A frozen, scale-aware **Visual Geometry Encoder** extracts multi-scale 3D priors from asynchronous RGB streams (10 Hz). These features are fused with high-frequency proprioceptive (50 Hz) via a **Multi-Head Cross-Attention mechanism**. To bridge the sim-to-real gap, a dual-head auxiliary decoder (middle right) regularizes the latent space by reconstructing local terrain topography and predicting system dynamics during training.

inefficiency and overfitting to simulation artifacts. Conversely, GeoLoco integrates a frozen VFM prior with physically grounded auxiliary regularization. This formulation explicitly circumvents representation collapse, enabling highly sample-efficient training and robust zero-shot deployment in unstructured real-world environments.

III. METHOD

This section presents GeoLoco, a novel geometry-aware humanoid locomotion policy that relies exclusively on a monocular RGB camera and proprioception.

A. Problem Formulation

We formulate the humanoid locomotion task as a Partially Observable Markov Decision Process (POMDP) solved via Proximal Policy Optimization (PPO) [15]. At each control step t , the actor policy π_θ generates the joint position commands $\mathbf{a}_t \in \mathbb{R}^{d_a}$. The observation space consists of the instantaneous proprioceptive state $\mathbf{o}_t \in \mathbb{R}^{d_p}$, the target velocity command $\mathbf{c}_t \in \mathbb{R}^3$, and the monocular RGB image $\mathbf{I}_t \in \mathbb{R}^{112 \times 112 \times 3}$. To mitigate the inherent scale ambiguity of monocular vision and transient sensory occlusions, the policy is explicitly conditioned on a proprioceptive history of length h , defined as an ordered sequence $\mathbf{o}_t^H = (\mathbf{o}_{t-h+1}, \dots, \mathbf{o}_t)$. The critic has access to privileged simulator state $\mathbf{s}_t \in \mathbb{R}^{371}$

to provide richer value estimates during training. Include ground-truth base linear velocity and height, foot-end kinematics, 3D contact forces, and a dense 17×11 local elevation map M_t .

To demonstrate the robustness of our RGB-driven visual representation, We employ a standard task-agnostic reward formulation prioritizing velocity tracking and joint smoothness, forcing the policy to rely strictly on 3D visual priors rather than reward engineering.

B. Geometry-Prior Visual Representation from RGB

To navigate complex terrains like stairs, a humanoid robot requires an awareness of the local 3D geometry that is invariant to lighting and texture. Instead of training a visual encoder from scratch, which often overfits to simulation-specific textures and fails in sim-to-real transfer, we leverage a pretrained VFM as a frozen 3D geometric prior. This paradigm shift allows the policy to operate on a high-dimensional representation that inherently encodes metric-relative depth and surface structures, effectively treating monocular RGB as a 3D latent surface rather than a flat 2D pixel array.

1) *Multi-Scale Tokenized Feature Extraction:* Given a monocular RGB frame $\mathbf{I}_t \in \mathbb{R}^{112 \times 112 \times 3}$, the frozen VFM tokenizes it into $N = 64$ non-overlapping patches, arranged

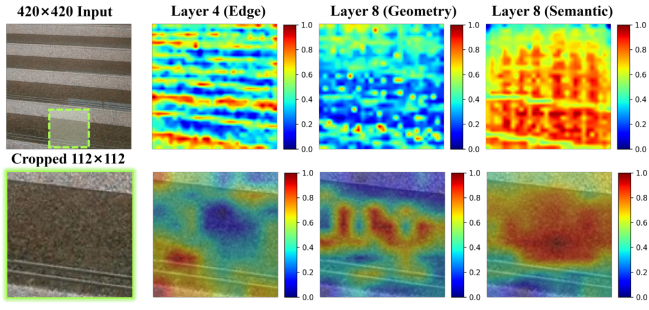


Fig. 3: Multi-Scale geometric activations from the frozen VFM. We visualize patch token activations from intermediate transformer layers (4/8/12).

in an 8×8 spatial grid. We instantiate the VFM using Depth-Anything-V2 (ViT-S) [8] with embedding dimension $d_v = 384$. Specifically, we leverage its metric-depth variant, which aligns perfectly with our objective of recovering absolute metric scale. Rather than using the scalar depth output, we extract the spatially-organized patch tokens from multiple intermediate transformer blocks ℓ .

Within the ViT-based architecture, shallow layers capture high-frequency geometric primitives, while deeper layers provide broader structural context [16], [8]. To exploit this hierarchy, we extract tokens from layers $\ell \in \{4, 8, 12\}$ of the 12-block ViT-S encoder, yielding three feature maps, each with $N = 64$ tokens of dimension $d_v = 384$.

2) *Channel-Grouped Spatial Projection*: Raw VFM embeddings occupy a high-dimensional semantic space unsuitable for high-frequency control. To adapt these representations into a compact spatial descriptor, we apply a parameter-free channel grouping operation independently to each layer. For each layer ℓ , the N patch tokens are first reshaped into a spatial feature map $\mathbf{F}_t^{(\ell)} \in \mathbb{R}^{d_v \times h \times w}$ where $h = w = 8$. We then partition the $d_v = 384$ channels into $G = 32$ groups of 12 and average within each group:

$$\mathbf{Z}_t^{(\ell)} = \text{GroupAvg}(\mathbf{F}_t^{(\ell)}) \in \mathbb{R}^{G \times h \times w}. \quad (1)$$

The three layers are then concatenated along the channel dimension to form the multi-scale geometric descriptor:

$$\mathbf{T}_t = \text{Concat}(\mathbf{Z}_t^{(4)}, \mathbf{Z}_t^{(8)}, \mathbf{Z}_t^{(12)}), \quad (2)$$

yielding $\mathbf{T}_t \in \mathbb{R}^{96 \times 8 \times 8}$, a compact spatial feature map where each of the $N = 64$ spatial locations carries a $d = 3G = 96$ dimensional descriptor. Compared to the raw concatenated ViT output, this represents an aggressive $12 \times$ per-token compression while strictly preserving the 8×8 spatial layout critical for terrain geometry. Crucially, this operation is entirely parameter-free and operates within the frozen encoder, introducing no additional learnable weights.

3) *Temporal Alignment and Asynchronous Inference*: A pivotal challenge in RGB-driven locomotion is the computational latency of large-scale VFMs. While the proprioceptive state \mathbf{o}_t is sampled at 50 Hz for reactive balancing, the VFM inference is decoupled and runs asynchronously at 10 Hz. Let Δt_{vis} denote the visual update interval in terms

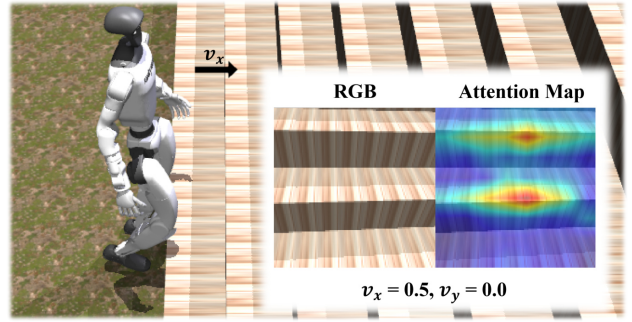


Fig. 4: Visualization of attention heatmap produced by the proprioceptive-query cross-attention module for a representative stair scene.

of control steps. For any high-frequency control step t , the policy consumes the most recently cached visual descriptor \mathbf{T}_τ , where the visual timestamp τ is updated via a zero-order hold mechanism:

$$\tau = \left\lfloor \frac{t}{\Delta t_{\text{vis}}} \right\rfloor \Delta t_{\text{vis}}. \quad (3)$$

This asynchronous design ensures that the humanoid’s reactive control loop remains unconstrained by the visual processing latency, while consistently benefiting from periodic 3D geometric corrections.

C. Multi-Head Cross-Attention Fusion

To bridge the gap between high-dimensional geometric representations and reactive motor control, we propose a Multi-Head Cross-Attention mechanism. Unlike conventional methods utilizing static convolution or naive concatenation, our approach treats humanoid locomotion as an active perception process, where the robot’s instantaneous physical state dynamically modulates its visual focus.

1) *Spatio-Temporal Token Formulation*: Instead of reasoning over a single frame, we aggregate geometry-prior descriptors from a short visual window of the most recent two asynchronous updates. The spatial feature map \mathbf{T}_τ from each frame f is flattened along its spatial dimensions to yield $N = 64$ sequence tokens, each of dimension $d = 96$. To enable the attention mechanism to distinguish tokens across time, we inject a learnable temporal embedding $\mathbf{e}_{\text{frame}}(f) \in \mathbb{R}^d$ into each token:

$$\tilde{\mathbf{t}}_{i,f} = \mathbf{t}_{i,f} + \mathbf{e}_{\text{frame}}(f), \quad i \in \{1, \dots, N\}, f \in \{0, 1\}. \quad (4)$$

These tokens are then concatenated to form a single spatio-temporal sequence $\tilde{\mathbf{T}}_\tau \in \mathbb{R}^{M \times d}$, where $M = 128$. This formulation empowers the subsequent attention mechanism to perform patch-level temporal alignment and deduce the relative motion of terrain features.

2) *Proprioceptive-Query Mechanism*: Our fusion formulation posits that the robot’s instantaneous kinematic and dynamic state must actively modulate spatial attention. Consequently, the instantaneous proprioceptive observation $\mathbf{o}_t \in \mathbb{R}^{d_o}$ ($d_o = 96$) is explicitly projected as the attention query \mathbf{q}_t , whereas the compressed proprioceptive history \mathbf{h}_t is decoupled to govern high-level policy execution and velocity estimation.

We project the current state into a query \mathbf{q}_t , and the enriched visual tokens into keys \mathbf{K}_τ and values \mathbf{V}_τ :

$$\begin{aligned} \mathbf{q}_t &= \mathbf{W}_Q \mathbf{o}_t \in \mathbb{R}^{1 \times d_a}, \\ \mathbf{K}_\tau &= \tilde{\mathbf{T}}_\tau \mathbf{W}_K \in \mathbb{R}^{M \times d_a}, \quad \mathbf{V}_\tau = \tilde{\mathbf{T}}_\tau \mathbf{W}_V \in \mathbb{R}^{M \times d_a}, \end{aligned} \quad (5)$$

where $d_a = 64$ is the attention embedding dimension, with $\mathbf{W}_Q \in \mathbb{R}^{d_o \times d_a}$ and $\mathbf{W}_K, \mathbf{W}_V \in \mathbb{R}^{d \times d_a}$. We employ multi-head attention (MHA) with 8 heads to allow parallel focus on multiple task-critical spatial regions. For the j -th attention head, the localized attention weights $\alpha_t^{(j)}$ are computed as:

$$\alpha_t^{(j)} = \text{softmax} \left(\frac{\mathbf{q}_t^{(j)} (\mathbf{K}_\tau^{(j)})^\top}{\sqrt{d_{\text{head}}}} \right) \in \mathbb{R}^{1 \times M}. \quad (6)$$

The outputs from all heads are concatenated and projected through a final linear layer with ELU activation, producing the fused visual context $\mathbf{c}_t^{\text{atten}} \in \mathbb{R}^{d_z}$.

3) *Physical Interpretation of the Attention Score*: Our formulation inherently grounds the attention weights α_t into three physical properties: semantic matching to extract traversable geometric primitives, spatial preference to bias visual focus based on postural inclination, and temporal preference to weigh historical versus instantaneous frames conditioned on locomotion velocity.

The final policy input is formed by concatenating the instantaneous state, the temporal proprioceptive history, and the state-dependent visual context: $\mathbf{u}_t^{\text{in}} = [\mathbf{o}_t; \mathbf{h}_t; \mathbf{c}_t^{\text{atten}}] \in \mathbb{R}^{288}$, ensuring a clean decoupling between reactive balance and geometry-aware path planning.

D. Geometry-Aware Auxiliary Estimator

To ensure that the fused latent space captures metric-accurate geometry rather than mere visual textures, we introduce an auxiliary representation learning scheme. During training, two lightweight decoding heads provide direct geometric supervision to the proprioceptive and visual pathways:

- **Velocity Estimation Head**: Estimates the robot’s base linear velocity $\hat{\mathbf{v}}_t = g_{\text{vel}}(\mathbf{h}_t) \in \mathbb{R}^3$ from the history latent alone. This provides an explicit supervisory gradient to the history encoder, complementary to the implicit RL policy gradient.
- **Terrain Reconstruction Head**: Reconstructs a front-facing subset of the local height map $\hat{\mathbf{M}}_t = g_{\text{recon}}(\mathbf{u}_t^{\text{in}}) \in \mathbb{R}^{99}$ from the full policy input. Only the forward half ($x \geq 0$) of the grid is predicted, focusing supervision on the terrain the robot is about to traverse.

The corresponding ground-truth signals \mathbf{v}_t and \mathbf{M}_t are obtained directly from the simulator. We define the total objective function as a combination of the PPO surrogate loss and the auxiliary regularizers:

$$\mathcal{L} = \mathcal{L}_{\text{PPO}} + \lambda_{\text{vel}} \text{MSE}(\hat{\mathbf{v}}_t, \mathbf{v}_t) + \lambda_{\text{recon}} \text{MSE}(\hat{\mathbf{M}}_t, \mathbf{M}_t), \quad (7)$$

where λ_{vel} and λ_{recon} are weighting coefficients. At deployment, the auxiliary heads are discarded, introducing zero computational overhead.

TABLE I: Visual Domain Randomization Parameters

Parameter	Randomization Range	Unit
Dome Light Intensity	[300, 2000]	lux
Directional Light Intensity	[1000, 6000]	lux
Light Color Variation	± 30	%
Lighting Angles	[0, 2π]	rad
Terrain Material	Random Texture Set	–
Camera Exposure/Contrast	[0.8, 1.2]	factor
Camera Position (x)	[−10, 10]	mm
Camera Position (y)	[−10, 10]	mm
Camera Position (z)	[−10, 10]	mm
Motion Blur Probability	15	%
Random Latency	[0, 100]	ms

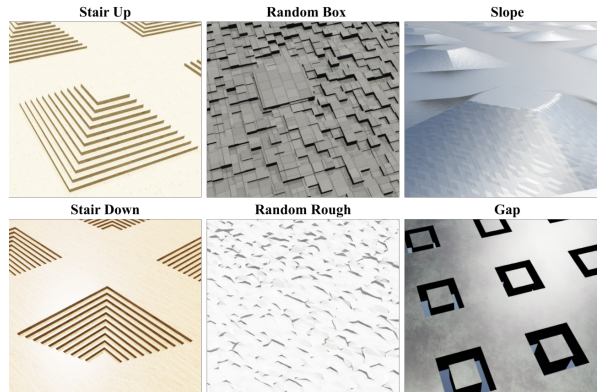


Fig. 5: Training environments featuring diverse geometries and randomized textures. This prevents overfitting to visual appearances, ensuring the policy learns robust 3D priors for sim-to-real transfer.

E. Sim-to-Real Training for RGB Policy

A primary hurdle in transitioning from simulation to real-world deployment for RGB-based policies is the significant discrepancy in visual appearance and hardware dynamics. To ensure the zero-shot transferability of GeoLoco, we employ a comprehensive domain randomization strategy encompassing both physical properties and visual modalities.

To force the policy to prioritize the underlying 3D geometric structures over superficial 2D textures, we apply extensive randomization to the visual environment in IsaacLab [1]. By varying lighting intensities, color temperatures, and shadow directions shown in Tab. I, we encourage the policy to extract lighting-invariant features. Furthermore, we randomly switch terrain materials during training, which prevents the agent from overfitting to specific simulation artifacts and ensures it relies on the metric-relative depth cues provided by the VFM. Critically, to account for the asynchronous 10 Hz visual update frequency against the 50 Hz control loop, we inject random latency [0, 100] ms into the visual feature stream.

IV. EXPERIMENTS

A. Experimental Setup

We compare GeoLoco with three categories of baselines: Proprioception-only DreamWaq [3]): Represents the lower bound without terrain awareness; Active Geometric Sensors (PIE [17], MORE [18], PLANC [19]): Use depth maps or

TABLE II: Benchmarked Comparison in Simulation. We compare GeoLoco against blind, depth-based, and other RGB-based baselines across various terrain difficulties. Best RGB-based result in bold.

Method	Modality	Stairs-Up \uparrow		Stairs Down		Slope		Gap	
		$R_{succ}(\%, \uparrow)$	$R_{trav}(\%, \uparrow)$	$R_{succ}(\%, \uparrow)$	$R_{trav}(\%, \uparrow)$	$R_{succ}(\%, \uparrow)$	$R_{trav}(\%, \uparrow)$	$R_{succ}(\%, \uparrow)$	$R_{trav}(\%, \uparrow)$
Medium Terrain Difficulty									
DreamWaq [3]	Blind	41.28	44.10	61.73	64.95	79.62	82.88	3.84	4.92
PIE [17]	Depth	90.63	92.10	93.22	94.58	98.91	99.42	92.44	93.76
MoRE [18]	Depth	81.94	83.21	84.10	85.67	96.72	97.88	82.35	83.90
PLANC [19]	Height map	91.37	92.76	94.05	95.44	99.36	99.72	93.18	94.92
CNN	RGB	59.42	60.88	70.35	71.90	84.72	85.96	29.58	30.74
GaussGym [20]	RGB	72.18	73.42	75.26	76.61	84.92	86.05	57.84	58.97
GeoLoco	RGB	82.76	84.09	85.33	86.62	97.10	98.25	83.05	84.66
Hard Terrain Difficulty									
DreamWaq [3]	Blind	11.24	13.06	29.87	32.41	58.92	62.17	0.00	0.00
PIE [17]	Depth	84.78	86.25	88.91	90.34	98.36	99.10	61.52	62.89
MoRE [18]	Depth	70.41	71.88	72.36	73.91	85.77	87.09	54.86	56.20
PLANC [19]	Height map	86.12	87.54	90.23	91.65	99.08	99.62	64.37	65.88
CNN	RGB	28.11	30.19	60.38	61.90	69.55	70.88	17.86	19.40
GaussGym [20]	RGB	49.42	51.88	70.15	71.62	74.38	75.91	24.86	25.72
GeoLoco	RGB	66.27	67.91	64.73	66.15	90.41	91.68	49.62	51.08




Fig. 6: **Real-world zero-shot deployment on Unitree G1.** The policy successfully traverses challenging geometric discontinuities under varied conditions, including (a) dark background stairs, (b) a gap obstacle, (c) high-riser stairs, and (d) a slope.

height fields, representing the current perceptive locomotion standard; and RGB-based End-to-End (CNN, GaussGym [20]). The training and simulation environments are implemented in IsaacLab [1]. Following standard perceptive locomotion benchmarks [6], we quantitatively evaluate the policies using two primary metrics:

- **Success Rate (R_{succ}):** The percentage of episodes where the humanoid successfully traverses the entire predefined length of the target terrain without triggering any termination conditions.
- **Traversal Rate (R_{trav}):** The average ratio of the forward distance traveled before an episode terminates to the total length of the terrain. This metric provides a soft penalty for partial successes.

Evaluation Protocol: During evaluation, we leverage massively parallel simulation, deploying 1,024 independent

environments for each terrain setting. Each environment operates continuously and resets automatically upon episode termination. This process iterates until 5 complete trials are collected per environment. Our curriculum scales difficulty from 0.0 to 1.0 (e.g., max stair height 0.25 m, max gap 0.3 m). To systematically probe the robustness boundaries of different modalities, we evaluate on two distinct tiers: *Medium* (level 0.4) and *Hard* (level 0.8).

Network Architecture and Implementation Details: For the visual perception pipeline, we adopt the small variant of the Depth-Anything-V2 (Metric Depth) [8] model as our frozen VFM backbone. The model utilizes a Vision Transformer (ViT-S) encoder [21] with approximately 24.8 M parameters. For the proprioceptive history, we set the window length to $h = 5$.

TABLE III: Quantitative ablation results in simulation. We report the maximum Mean Difficulty Level reached through curriculum learning and the corresponding Success Rate $R_{\text{succ}}\%$ on the final curriculum stage. All results are averaged over 5 random seeds.

Category	Variant	Modality	Fusion Strategy	Difficulty \uparrow	$R_{\text{succ}}(\%, \uparrow)$
Full Method	GeoLoco (Ours)	RGB	Cross-Attention	0.87	86.4
(A) Representation	Scratch CNN (End-to-End)	RGB	Cross-Attention	0.58	60.4
	Semantic VFM (DINOv2)	RGB	Cross-Attention	0.61	66.1
	CLS Tokens	RGB	Cross-Attention	0.55	58.7
	Single-Layer Patch Token	RGB	Cross-Attention	0.75	74.8
(B) Modality Fusion	Naive MLP Concatenation	RGB	Concatenation	0.76	71.2
	w/o Proprioceptive History (h_t)	RGB	Cross-Attention	0.64	67.2
	Single-frame Vision (10 Hz)	RGB	Concatenation	0.77	79.5
	Temporal Vision (10 Hz)	RGB	Conv2d	0.79	83.5
(C) Regularization	w/o Terrain Reconstruction	RGB	Cross-Attention	0.71	74.2
	w/o Velocity Estimation	RGB	Cross-Attention	0.81	81.4
	w/o Visual Domain Randomization	RGB	Cross-Attention	0.80	84.1
(D) Baselines	Proprioceptive Only (Blind)	–	None	0.31	20.1
	Depth-based baseline [18]	Depth	Concatenation	0.94	93.2
	Height maps [5]	LiDAR	Concatenation	0.95	94.5

B. Simulation Results

We report the quantitative evaluation across varying terrain difficulties in Tab. II. The results reveal several critical insights regarding sensory modalities and visual representation learning in perceptive locomotion. First, the proprioceptive-only baseline [3] completely fails on severe geometric discontinuities, yielding a 0.00% success rate R_{succ} on Hard Gaps. This confirms the absolute necessity of exteroception for negotiating non-continuous terrains.

Second, standard RGB-driven policies exhibit severe vulnerability to environmental complexity. While end-to-end CNN and GaussGym [20] demonstrate moderate success on Medium terrains, their performance degrades sharply under Hard conditions. In contrast, GeoLoco substantially mitigates this degradation, maintaining a robust R_{succ} of 66.27% on Hard Stairs-Up and 49.62% on Hard Gaps. This substantial margin validates that conceptualizing monocular RGB as a 3D latent representation explicitly recovers the metric-level geometric cues inherently lacking in raw 2D pixel observations.

Finally, we benchmark GeoLoco against depth and heightmap-based methods, which utilize explicit geometric sensors. Remarkably, our purely RGB-driven framework achieves performance competitive with the depth-based MoRE baseline [18]. On Medium Stairs-Up, GeoLoco (82.76%) even marginally exceeds MoRE (81.94%). A predictable performance gap remains between our method and explicit heightmap algorithms on Hard terrains due to the physical limitations of monocular scale ambiguity.

C. Real-World Deployment on Unitree G1

We deploy the trained GeoLoco policy zero-shot on the 29-DoF Unitree G1 humanoid. The control loop operates at 50 Hz, mirroring the simulation environment. Exteroception relies solely on the RGB stream of an onboard Intel RealSense D435i camera, with policy inference executed on a desktop RTX 4090 GPU.

TABLE IV: Quantitative Real-World Performance on Unitree G1. Each method was evaluated 10 trials per terrain scenario. Best result in bold.

Method	Stairs (0.23 m riser)		Gap (0.25 m width)	
	$R_{\text{succ}}(\%, \uparrow)$	Time (\downarrow)	$R_{\text{succ}}(\%, \uparrow)$	Time (\downarrow)
Proprioceptive (Blind)	30%	10.1	0%	6.5
CNN-based	40%	8.2	40%	5.8
GeoLoco (Ours)	80%	5.4	70%	4.2

Quantitative Evaluation. To rigorously assess real-world reliability, we benchmark GeoLoco against both blind and end-to-end CNN baselines across 10 consecutive trials per terrain, as summarized in Tab. IV. GeoLoco achieves success rates of 80% and 70% on 0.23 m stairs and 0.25 m gaps, respectively, substantially outperforming the baselines. Furthermore, the VFM-derived geometric priors enable proactive traversals, consistently reducing the average execution time to 5.4 s on stairs and 4.2 s on gaps.

Qualitative Dynamic Behaviors. Qualitatively, the physical deployment validates several robust dynamic capabilities. On severe geometric discontinuities like steep stairs and wide gaps shown in Fig. 6b, c, the robot refrains from naive velocity tracking, exhibiting instead anticipatory leg-lifting and adaptive foot placement. Furthermore, GeoLoco successfully negotiates stairs under low-light conditions (Fig. 6a), confirming that the extracted 3D latent representations inherently resist superficial illumination shifts. Finally, when traversing continuous slopes (Fig. 6d), the policy dynamically regulates the torso pitch and center of mass, achieving geometry-aware whole-body coordination without explicit terrain-specific reward engineering.

D. Ablation Analysis

Table III presents extensive ablations to isolate the contributions of three core components:

1) *Efficacy of 3D Geometric Priors:* As reported in Tab. III(A), replacing the frozen geometric encoder with an end-to-end trained CNN precipitates a severe performance

degradation, reducing the success rate R_{succ} from 86.4% to 60.4%. Furthermore, employing a state-of-the-art semantic foundation model DINOv2 [16] yields marginal benefits, achieving an R_{succ} of only 66.1%. This substantial performance gap substantiates that while semantic models excel at object-level recognition, they fundamentally lack the precise metric-scale spatial grounding required for negotiating discrete geometric discontinuities.

2) *Impact of Multi-Head Cross-Attention Fusion:* Table III(B) validates that our proposed fusion strategy is critical for managing high-dimensional VFM priors. Reverting to MLP Concatenation incurs a 15.2% absolute drop in success rate. This degradation suggests that unconstrained feature concatenation renders the policy susceptible to overfitting transient visual artifacts. Conversely, the Proprioceptive-Query Cross-Attention mechanism effectively modulates the visual focus toward terrain features conditioned on the instantaneous state. Furthermore, ablating the Proprioceptive History \mathbf{h}_t severely diminishes the R_{succ} to 67.2%, highlighting the indispensability of temporal context for disambiguating partial visual occlusions induced by the robot's own limbs.

3) *Regularization as Physical Grounding:* The dual-head auxiliary estimator functions as an explicit regularization mechanism to inherently ground the high-dimensional latent space in physical geometry. Empirically, we observed that all auxiliary losses stably converge within 2,000 steps across diverse visual representations. Ablating the Terrain Reconstruction head induces a substantial performance decline, decreasing the R_{succ} from 86.4% to 74.2%. This validates that auxiliary geometric supervision is imperative for preventing representation collapse and ensuring robust terrain traversal.

V. CONCLUSION

We present GeoLoco, a purely RGB-driven reinforcement learning framework for humanoid locomotion. By conceptualizing monocular RGB as high-dimensional 3D latent representations via a frozen Visual Foundation Model (VFM), our approach achieves efficient visual-proprioceptive fusion through a lightweight cross-attention mechanism. Extensive simulations demonstrate that our dual-head auxiliary learning is crucial for explicitly grounding these visual priors in physical geometry. In zero-shot real-world deployments on the Unitree G1, the policy successfully negotiates challenging terrains such as stairs and low-light environments. These qualitative results validate the feasibility and robustness of VFM-derived features against visual degradation. Ultimately, GeoLoco provides a scalable paradigm that bridges geometric perception and semantic context, paving the way for next-generation Vision-Language-Action (VLA) agents.

REFERENCES

[1] M. Mittal, C. Yu, Q. Yu, J. Liu, N. Rudin, D. Hoeller, J. L. Yuan, R. Singh, Y. Guo, H. Mazhar, A. Mandlekar, B. Babich, G. State, M. Hutter, and A. Garg, "Orbit: A unified simulation framework for interactive robot learning environments," *IEEE Robotics and Automation Letters*, vol. 8, no. 6, pp. 3740–3747, 2023.

[2] 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," *arXiv preprint arXiv:2108.10470*, 2021.

[3] I. M. Aswin Nahrendra, B. Yu, and H. Myung, "Dreamwaq: Learning robust quadrupedal locomotion with implicit terrain imagination via deep reinforcement learning," in *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pp. 5078–5084, 2023.

[4] W. Cui, S. Li, H. Huang, B. Qin, T. Zhang, J. Han, L. Zheng, Z. Tang, C. Hu, N. Yan, J. Chen, and Z. Jiang, "Adapting humanoid locomotion over challenging terrain via two-phase training," in *Proc. of the Conf. on Robot Learning (CoRL)*, 2024.

[5] Z. Zhuang, S. Yao, and H. Zhao, "Humanoid parkour learning," *arXiv preprint arXiv:2406.10759*, 2024.

[6] H. Wang, Z. Wang, J. Ren, Q. Ben, T. Huang, W. Zhang, and J. Pang, "Beamdojo: Learning agile humanoid locomotion on sparse footholds," in *Proc. of Robotics: Science and Systems (RSS)*, 2025.

[7] X. Cheng, K. Shi, A. Agarwal, and D. Pathak, "Extreme parkour with legged robots," in *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2024.

[8] L. Yang, B. Kang, Z. Huang, Z. Zhao, X. Xu, J. Feng, and H. Zhao, "Depth anything v2," in *Proc. of the Advances in Neural Information Processing Systems (NIPS)*, 2024.

[9] H. Xue, X. Huang, D. Niu, Q. Liao, T. Kragerud, J. T. Gravdahl, X. B. Peng, G. Shi, T. Darrell, K. Sreenath, and S. Sastry, "Leverb: Humanoid whole-body control with latent vision-language instruction," *arXiv preprint arXiv:2506.13751*, 2025.

[10] T. Lin, G. Li, Y. Zhong, Y. Zou, and B. Zhao, "Evo-0: Vision-language-action model with implicit spatial understanding," *arXiv preprint arXiv:2507.00416*, 2025.

[11] S. Zeng, D. Qi, X. Chang, F. Xiong, S. Xie, X. Wu, S. Liang, M. Xu, and X. Wei, "Janusvln: Decoupling semantics and spatiality with dual implicit memory for vision-language navigation," in *Proc. of the Intl. Conf. on Learning Representations (ICLR)*, 2026.

[12] T. He, Z. Wang, H. Xue, Q. Ben, Z. Luo, W. Xiao, Y. Yuan, X. Da, F. Castañeda, S. Sastry, C. Liu, G. Shi, L. Fan, and Y. Zhu, "Viral: Visual sim-to-real at scale for humanoid loco-manipulation," *arXiv preprint arXiv:2511.15200*, 2025.

[13] H. Xue, T. He, Z. Wang, Q. Ben, W. Xiao, Z. Luo, X. Da, F. Castañeda, G. Shi, S. Sastry, L. J. Fan, and Y. Zhu, "Opening the sim-to-real door for humanoid pixel-to-action policy transfer," *arXiv preprint arXiv:2512.01061*, 2025.

[14] Z. Luo, C. Tessler, T. Lin, Y. Yuan, T. He, W. Xiao, Y. Guo, G. Chechik, K. Kitani, L. Fan, and Y. Zhu, "Emergent active perception and dexterity of simulated humanoids from visual reinforcement learning," *arXiv preprint arXiv:2505.12278*, 2025.

[15] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.

[16] M. Oquab, T. Darcet, T. Moutakanni, H. V. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. Haziza, F. Massa, A. El-Nouby, R. Howes, P.-Y. Huang, H. Xu, V. Sharma, S.-W. Li, W. Galuba, M. Rabbat, M. Assran, N. Ballas, G. Synnaeve, I. Misra, H. Jegou, J. Mairal, P. Labatut, A. Joulin, and P. Bojanowski, "Dinov2: Learning robust visual features without supervision," in *Proc. of the IEEE/CVF Intl. Conf. on Computer Vision (ICCV)*, 2023.

[17] 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 (RA-L)*, pp. 9986–9993, 2024.

[18] 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.

[19] M. Dai, W. D. Compton, J. Li, L. Yang, and A. D. Ames, "Walk the planc: Physics-guided rl for agile humanoid locomotion on constrained footholds," *arXiv preprint arXiv:2601.06286*, 2026.

[20] A. Escontrela, J. Kerr, A. Allshire, J. Frey, R. Duan, C. Sferazza, and P. Abbeel, "Gaussgym: An open-source real-to-sim framework for learning locomotion from pixels," *arXiv preprint arXiv:2510.15352*, 2025.

[21] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," in *Proc. of the Intl. Conf. on Learning Representations (ICLR)*, 2021.