SURE: SUPPORT-CONDITIONED UNRELIABILITY-AWARE RESIDUAL EXTEROCEPTION FOR GRACEFUL DEGRADATION IN HUMANOID LOCOMOTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Visual terrain perception helps humanoids traverse complex terrain, but erroneous or stale elevation maps can misdirect foot placement and whole-body motion. Controllers therefore need to determine when to stop relying on map guidance while maintaining basic locomotion after a mismatch. We propose SURE (Support-Conditioned Unreliability-Aware Residual Exteroception), a visual-residual control method based on physical consistency between vision and proprioception. SURE adds a visual residual to a blind base policy driven by proprioception, providing terrain-dependent action corrections. To decide when to use these corrections, it infers foot support and collision states from proprioception and checks them against local map geometry. Vertical checks compare supporting-foot height with mapped terrain height, while horizontal checks compare foot-front obstruction with mapped obstacle geometry. Physical contradictions prompt residual withdrawal, leaving the blind base policy in control. Subsequent foot placements must pass support verification before visual corrections resume. Simulations show that SURE preserves complex-terrain locomotion under nominal vision. Under local geometric mismatch and map delay, SURE improves walking success relative to the same policy with the visual residual always enabled. Restoration ablations show that map-based support verification improves completion beyond proprioceptive recontact under delay and height drift. Real-robot experiments further demonstrate terrain traversal with omitted obstacles or delayed maps. These results show that physically grounded residual control reconciles terrain adaptation under nominal vision with basic locomotion after perceptual mismatch, providing interpretable graceful degradation.

1 INTRODUCTION

Visual terrain preview gives humanoid robots geometric information before contact, enabling anticipatory adjustments to foot placement and whole-body motion. This information supports locomotion over stairs, gaps, and sparse footholds (Zhuang et al., 2024; Long et al., 2025; Wang et al., 2025b; Sun et al., 2025; 2026; Zhu et al., 2026). However, occlusion and missed scans can hide terrain protrusions and depressions. Delayed updates and imperfect motion compensation can also misalign the map with the terrain. These errors may occur without producing any explicit fault alarm. A controller may therefore act on incorrect support heights or obstacle locations, causing misplaced steps, unexpected collisions, or loss of balance (Figure 1).

Existing approaches use multimodal fusion and training with noisy observations to improve tolerance to perception errors (Miki et al., 2022; Zhang et al., 2025a). Another approach combines vision and blind policies, using expected-return estimates to select between them when perception becomes unreliable (Ren et al., 2025). We focus on the physical basis for fallback while a map continues to arrive. The question is whether foot–terrain interaction can reveal when the map no longer provides suitable guidance. Our objective is *graceful degradation*, prioritizing basic proprioception-driven locomotion under perceptual mismatch even at the cost of reduced visual terrain adaptation.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

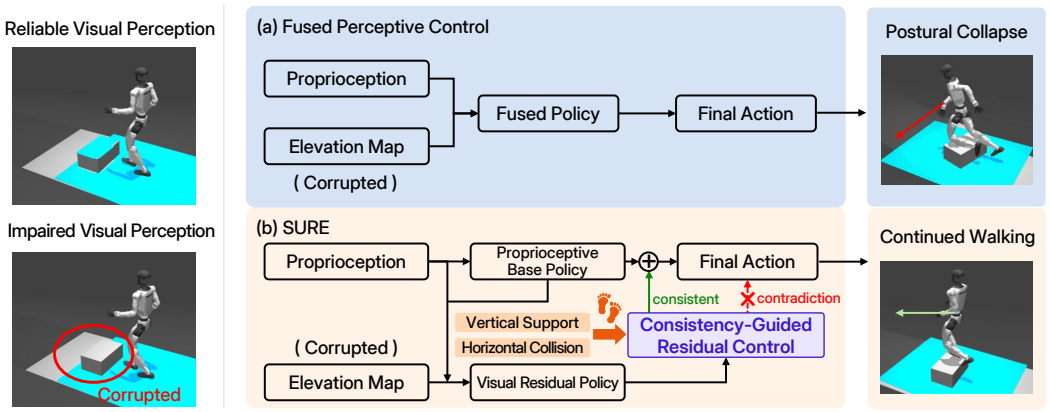


Figure 1: Motivation for physically grounded visual fallback. Incorrect elevation maps can mislead fused perceptive control. SURE makes visual corrections removable while retaining a blind base policy. Physical contradictions guide withdrawal, and subsequent support verification guides restoration.

We propose SURE (Support-Conditioned Unreliability-Aware Residual Exteroception), which uses local physical consistency between vision and proprioception to govern visual corrections. A blind base policy continually responds to proprioceptive feedback, while a visual residual adds terrain-dependent corrections. This decomposition makes visual corrections independently removable without interrupting the base policy. Later foot placements that agree with mapped support surfaces provide evidence for restoring visual corrections.

SURE estimates each foot's support and collision states from proprioception alone and compares them with map geometry through two checks. Vertical support consistency compares the supporting sole with mapped terrain height to identify local support-height mismatch. Horizontal collision consistency compares obstruction at the foot front with mapped obstacle geometry. Each consistency-triggered withdrawal can be traced to a particular foot and contradiction type, giving fallback an interpretable physical basis.

We evaluate nominal and faulty conditions through simulation comparisons, controlled ablations, and real-robot experiments. Our contributions are as follows.

- A visual-residual control mechanism that adds terrain-dependent corrections to an independently usable blind base policy, allowing visual corrections to be withdrawn separately.
- Physically distinct vertical support and horizontal collision checks that give residual withdrawal an interpretable, foot-specific physical basis.
- Simulation comparisons showing preserved nominal capability and higher success under map mismatch than the same-weight always-visual endpoint. Ablations separate withdrawal cues and establish the restoration benefit of map-based support verification beyond proprioceptive recontact.
- Real-humanoid terrain traversal with measured elevation maps and proprioceptive feedback under omitted obstacles or delayed maps, demonstrating consistency-driven residual withdrawal and restoration.

2 Related Work

Basic locomotion and visual terrain perception. Reinforcement learning supports sim-to-real legged control and adaptation (Hwangbo et al., 2019; Gu et al., 2024a; Radosavovic et al., 2024; Kumar et al., 2021; Cha et al., 2025). Related work covers blind stair traversal, locomotion through constrained terrain, and disturbance rejection (Siekmann et al., 2021; Xie et al., 2025; Zhang et al., 2025b). Depth images and elevation maps add terrain preview for anticipatory foot placement and whole-body motion (Agarwal et al., 2023; Zhuang et al., 2023; 2024; Wang et al., 2025b; Zhu et al.,

2026). Further work improves the extraction and representation of terrain geometry for perceptive locomotion (He et al., 2025; Sun et al., 2025; 2026; Song et al., 2026; Long et al., 2025).

Robust locomotion under unreliable perception. Sensor errors and localization uncertainty affect the reliability of terrain preview. Probabilistic elevation maps explicitly model sensor and localization uncertainty (Fankhauser et al., 2014; 2018). Related navigation methods combine terrain estimates with locomotion stability to guide traversability decisions (Muenprasitivej et al., 2024; Yoon et al., 2026). At the policy level, the aim is to extract useful control information from unreliable observations. Denoising world models improve proprioceptive state estimation (Gu et al., 2024b). Implicit–explicit representations and teacher–student distillation help policies handle unreliable visual inputs (Luo et al., 2024; Zhang et al., 2025a). Miki et al. use an attention-based recurrent encoder to combine complementary exteroceptive and proprioceptive information under unreliable observations (Miki et al., 2022).

Vision–blind composition and residual control. At the control level, fallback requires both a criterion for withdrawing visual guidance and a locomotion capability that does not depend on vision. VB-Com estimates the expected returns of vision and blind policies from proprioceptive history and uses them to select between the two policies (Ren et al., 2025). Residual reinforcement learning instead augments an existing controller with a learned correction (Johannink et al., 2019). SURE uses residual decomposition to keep a blind base policy active, retaining or withdrawing only the visual residual rather than selecting between two complete policies. Within this structure, it compares map geometry with foot interaction states inferred from proprioception. Physical consistency, rather than expected policy return, determines whether visual corrections remain active.

3 Problem Formulation

Perceptive locomotion. We consider a humanoid that follows motion commands over complex terrain using proprioception and a local elevation map. We model the task as a partially observable Markov decision process (POMDP). At time t, the observation is $\mathbf{o}_t = [\mathbf{o}_t^P, \widetilde{H}_t]$. Here $\mathbf{o}_t^P$ includes proprioceptive history, known motion commands, and action history. The map $\widetilde{H}_t$ is the robot-centric elevation map actually received by the controller. The policy selects $\mathbf{a}_t \sim \pi_\theta(\cdot \mid \mathbf{o}_t)$, where θ denotes its trainable parameters. The action $\mathbf{a}_t \in \mathbb{R}^D$ specifies normalized position offsets for D controlled joints. These offsets are converted to joint-position targets and tracked by a proportional–derivative (PD) controller. Appendix A.2 gives the execution interface. The learning objective is to maximize expected discounted return,

$$\max_\theta J(\pi_\theta), \qquad J(\pi_\theta) = \mathbb{E}_{\tau\sim\pi_\theta}\left[\sum_{t=0}^{\infty} \gamma^t r(\mathbf{s}_t, \mathbf{a}_t)\right]. \tag{1}$$

Here $\mathbf{s}_t$ is the environment state, r the task reward, and $\gamma \in [0, 1)$ the discount factor. The trajectory τ is generated by policy–environment interaction.

Control under map mismatch. At deployment, the received map may misrepresent support-surface heights or local obstacle geometry. We call this discrepancy a local map mismatch and focus on cases without an explicit fault alarm. The controller must use proprioceptive feedback and the received map to decide whether map guidance remains appropriate. Under nominal sensing, it should retain the locomotion capabilities afforded by terrain preview. Once a mismatch becomes observable, the objective is to prioritize basic proprioception-driven locomotion while allowing the benefits of visual terrain adaptation to diminish.

4 SURE

4.1 Visual residual control

Humans can continue walking with their eyes closed in controlled settings (Oliveira et al., 2017). This observation motivates separating basic locomotion from visually guided terrain adaptation. SURE

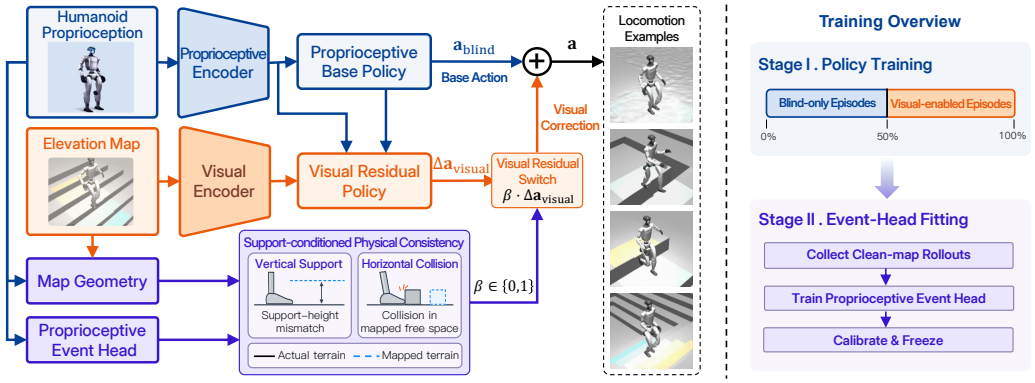


Figure 2: SURE architecture and training. A binary switch controls the visual residual while the blind base policy remains active. The proprioceptive event head and map geometry provide vertical support and horizontal collision checks, illustrated by two representative contradictions. After policy training, the event head is trained and calibrated on rollouts without injected visual faults.

therefore adds a visual residual to a blind base policy driven solely by proprioception (Figure 2). Omitting time subscripts in this section, the action at each control step is

$$\mathbf{a} = \mathbf{a}_{\text{blind}} + \beta \Delta \mathbf{a}_{\text{visual}}, \qquad \beta \in \{0, 1\}. \tag{2}$$

The blind branch produces the base action $\mathbf{a}_{\text{blind}}$ from proprioceptive history. The visual branch combines proprioceptive features, elevation-map features, and the base action to produce the terrain-dependent correction $\Delta \mathbf{a}_{\text{visual}}$. The global switch β enables the residual at one and withdraws it at zero. Withdrawal removes only the visually guided correction, while the blind base policy continues responding to current proprioception.

This structure makes visual corrections removable, but deciding when to withdraw them requires a criterion. SURE checks mapped terrain against the robot's own support and collision feedback. A contradiction between foot–terrain interaction and map geometry provides a physical reason to withdraw map-guided corrections.

4.2 Support-conditioned dual-axis physical consistency

Consistency checks compare each foot's interaction state with map geometry at the same location. The proprioceptive event head estimates per-foot support and collision-state probabilities from proprioceptive history without map input. Calibration and thresholding convert these probabilities into binary states. For foot $f \in \{L, R\}$, s_f and e_f denote inferred support and collision, with one indicating the corresponding event.

To compare these states with the corresponding map geometry, forward kinematics (FK) computes base-frame foot poses from the current joint angles. We then transform foot sampling points into the gravity-aligned local map frame. Each check requires valid event and geometric information. Invalid inputs do not produce a contradiction.

Vertical consistency during support. During the support phase, sole position provides a physical reference for support-surface height. During swing, however, a sole above the ground is normal. The vertical check therefore applies only to feet that the event head identifies as supporting.

Terrain such as a stair edge may support only part of the sole. If at least one sole sample agrees with the mapped height, the map can still explain this partial support. A vertical contradiction therefore requires every sole sample to differ from the mapped height by more than the tolerance. Let $z_{f,j}$ be the height of sole point j on foot f, and $H_{f,j}$ the received map height at that location. With height tolerance ϵ_s, the contradiction indicator is

$$c_f^v = s_f \, \mathbf{1}\left[\min_j |z_{f,j} - H_{f,j}| > \epsilon_s\right]. \tag{3}$$

Here $\mathbf{1}[\cdot]$ is one when the condition holds and zero otherwise. For example, support below the mapped surface can expose an unmapped depression.

Horizontal collision consistency. Whereas the vertical check compares support heights, the horizontal check compares foot-front obstruction with mapped obstacle geometry. An elevation map represents terrain through surface height at each horizontal location. Comparing foot-front samples with this surface thus distinguishes mapped free space from occupied terrain. Here j indexes foot-front samples, and $z_{f,j}$ and $H_{f,j}$ retain their meanings as sample height and corresponding map height. With a fixed height margin ϵ_g, define the map geometry state

$$g_f = \begin{cases} 0, & \forall j, \quad z_{f,j} > H_{f,j} + \epsilon_g, \\ 1, & \exists j, \quad z_{f,j} < H_{f,j} - \epsilon_g, \\ \perp, & \text{otherwise.} \end{cases} \tag{4}$$

Thus, $g_f = 0$ indicates that all front samples lie in mapped free space, while $g_f = 1$ indicates that at least one lies inside mapped terrain. The indeterminate state $g_f = \perp$ produces no horizontal contradiction.

The first contradiction occurs when the map indicates free space but proprioception indicates a collision, so $g_f = 0$ and $e_f = 1$. This comparison applies regardless of the inferred support state of that foot.

The second occurs when the foot front lies inside mapped terrain without an inferred collision, so $g_f = 1$ and $e_f = 0$. Normal support contact does not require a horizontal collision signal, so this comparison applies only to non-support feet ($s_f = 0$).

For $g_f \in \{0, 1\}$, the horizontal contradiction indicator is therefore

$$c_f^h = (1 - g_f)e_f + g_f(1 - s_f)(1 - e_f). \tag{5}$$

The two terms correspond to these two types of contradiction.

The vertical and horizontal checks use support height and obstruction feedback, respectively, and are applied separately to each foot. Either $c_f^v = 1$ or $c_f^h = 1$ can request withdrawal of the visual residual. Appendix A.3 specifies sampling, geometric tolerances, and validity conditions.

4.3 Consistency-driven residual control and training

The controller combines per-foot requests to govern a single residual switch β, but withdrawal and restoration use different criteria. A contradiction justifies withdrawal, but its disappearance alone does not establish renewed agreement between map geometry and physical interaction. Recontact confirms physical support, not that the map generating visual corrections describes that support correctly. Restoration therefore requires map-based support verification at subsequent foot placements. While observations remain valid, each foot must establish new support consistent with the map. Once withdrawal conditions have cleared and both feet have passed verification, continuous confirmation is required before re-enabling the residual. This verification is limited to the local footholds checked before restoration. Appendix A.3 gives the state updates and validity handling.

To keep the blind base policy independently usable, policy training uses a staged proximal policy optimization (PPO) curriculum. The first half of training uses only the blind base policy. During the second half, the proportion of vision-enabled episodes increases gradually until the blind-to-vision sampling ratio reaches $1 : 9$. Training controls visual availability per episode, whereas deployment controls β through physical consistency.

The visual branch applies stop-gradient to its proprioceptive-feature and base-action inputs, preventing backpropagation into the blind representation through these paths. The blind branch nevertheless continues learning through the combined action. Appendix A.2 details the networks, gradient paths, and training settings.

After policy training, we collect walking trajectories without injected visual faults and use simulated interactions to label per-foot support contact and current collision state. These labels supervise separate event-head training and calibration. Both training and deployment inputs comprise robot-available proprioception, derived FK features, and known commands and action history, without

simulator ground truth. Event-head parameters remain fixed at deployment, so consistency checks require no online fitting. Appendix C.1 describes event-head training and calibration.

5 EXPERIMENTS

We evaluate nominal perceptive locomotion (E1), control under map mismatch (E2), the physical cues for withdrawal and restoration (E3), and real-robot deployment (E4).

5.1 EVALUATION SETUP

We compare seven configurations. SURE and its fixed $\beta = 0$ and $\beta = 1$ configurations share policy weights, isolating dynamic withdrawal from the always-blind and always-visual endpoints. System-level comparisons include independently trained Blind, Visual, Heavy-DR Visual, and full-policy hard switching (HS). HS combines independent Blind and Visual policies with a matched proprioceptive event head and the same family of geometric and control rules. Appendix A.4 gives the configurations.

Clean comprises nine conditions, including rough ground, randomly placed low obstacles, gutters, and ascent and descent on slopes, high steps, and stairs.

We further evaluate five map faults with different consequences for foot–terrain interaction. In **G1, unmapped protrusion**, the map represents a real obstacle as flat ground, so foot-front collision feedback contradicts the mapped free space. Such omissions can arise from occlusion or weak returns from transparent or low-reflectance obstacles. Conversely, **G2, false support**, depicts a real depression as a continuous plane, so the indicated surface disagrees with actual foot support. This can occur when the ground subsides without a corresponding map update. In **G3, stale stair maps**, frames are delayed by 0.1–1.0 s, with the delay fixed within each episode. As the robot moves, stale stair edges become misaligned with the terrain encountered by the feet. In **G4, relative map-height drift**, maps continue to update, but the mapped surface gradually shifts below the unchanged flat ground. This creates a map-to-body height-registration error, making mapped height inconsistent with the physical support height. G1–G4 provide no explicit fault alarm, whereas **G5, map-stream interruption**, stops updates during walking and provides a timeout flag to SURE and HS. It therefore tests fallback after a reported sensor-stream fault, while fixed endpoints do not respond to the flag. Appendix A.5 specifies geometry and timing.

All methods share the walking command and termination criteria, with policy weights fixed throughout evaluation. Success requires reaching the finish within the time limit. Body collisions involving links other than the feet and ankles trigger termination and count as failures. Torso-tilt CVaR95 measures large posture deviations, while velocity-tracking error measures command following. Action-change rate and torso angular velocity describe how rapidly control commands and torso orientation change, respectively. Motion metrics include valid trajectories up to failure and are interpreted alongside success. Appendix A.6 specifies the criteria, aggregation, and sample sizes.

5.2 COMPARISONS AND MECHANISM ANALYSIS

E1. Preserving nominal perceptive locomotion. SURE achieved 99.93% nominal success, close to its always-visual endpoint at 99.99% and Visual at 99.96% (Table 1). The visual residual was particularly important on narrow repeated gutters. SURE achieved 99.33% success, compared with 0.37% for its blind endpoint and 0.83% for independent Blind (Appendix Table 8). The residual enabled traversal of this terrain where blind control largely failed, and consistency checks preserved that ability. The mean episode-level residual-on fraction was 96.16% under nominal sensing (Appendix Table 10). SURE therefore preserved visual corrections rather than completing nominal tasks through continuous blind control.

E2. Control under map faults. The same-weight endpoints reveal why neither always-visual nor always-blind control performs well across both nominal and faulty conditions. Always-blind control reached 100% success in every fault condition, G1–G5, but only 77.82% in Clean. Always-visual control achieved 99.99% in Clean but degraded under all five faults, with success ranging from 0% to 79.27%. SURE retained 99.93% in Clean, reached 100% in G1, G2, G4, and G5, and achieved

Table 1: Task completion and torso posture under nominal sensing and five map faults. SR is success rate (%) and T is torso-tilt CVaR95 (rad). Bold marks the best displayed value per metric and condition, including ties.

Configuration	Clean	G1	G2	G3	G4	G5
(a) Success rate (%) ↑						
SURE	99.93	**100.00**	**100.00**	95.87	**100.00**	**100.00**
Blind	69.01	**100.00**	**100.00**	**100.00**	**100.00**	**100.00**
Visual	99.96	0.00	0.00	10.03	0.00	0.00
Heavy-DR Visual	87.83	0.00	7.33	96.33	**100.00**	0.00
HS	97.72	99.67	99.90	40.63	**100.00**	**100.00**
SURE $\beta = 0$	77.82	**100.00**	**100.00**	**100.00**	**100.00**	**100.00**
SURE $\beta = 1$	**99.99**	79.27	28.67	10.00	0.10	0.00
(b) Torso tilt (rad) ↓						
SURE	0.0733	0.1071	0.0761	0.0698	0.0628	0.0618
Blind	**0.0579**	0.0887	0.0894	0.0675	**0.0449**	0.0676
Visual	0.0854	0.0811	0.1013	0.1221	0.3007	0.2729
Heavy-DR Visual	0.0927	**0.0794**	0.0955	0.0705	0.0642	0.0754
HS	0.0854	0.1115	0.0878	0.1095	0.0939	0.0771
SURE $\beta = 0$	0.0789	0.1119	**0.0639**	**0.0592**	0.0565	**0.0592**
SURE $\beta = 1$	0.0753	0.1548	0.1339	0.0803	0.3853	0.1026

95.87% in G3 (Table 1). Selective residual withdrawal thus reconciled nominal terrain adaptation with basic locomotion under map faults.

At the same 100% success rate in G1 and G2, SURE reduced velocity-tracking error relative to its blind endpoint by approximately 30.6% and 26.2% (Appendix Table 9). The mean residual-on fractions remained 82.66% and 72.88% (Appendix Table 10). These localized faults therefore did not require abandoning visual control throughout the traversal. The tracking gains are consistent with retaining useful visual corrections outside conflicting interactions and withdrawing them when physical feedback exposes map mismatch.

The independent baselines further distinguish robustness to one fault from performance across conditions. Heavy-DR Visual slightly exceeded SURE in G3 at 96.33%, but achieved 87.83%, 0%, and 7.33% in Clean, G1, and G2, respectively. Its tolerance of delayed maps did not extend to the other two geometric faults, whereas SURE preserved both nominal capability and high completion across these mismatches. Under G4, Visual achieved 0% success, while Heavy-DR Visual completed every task but had 3.66 times SURE's velocity-tracking error (Appendix Table 9). SURE's benefit in this condition therefore extended beyond task completion to command following.

In G1–G3, SURE achieved higher success and lower torso-tilt CVaR95 than both its always-visual endpoint and HS (Table 1). In G3, tilt decreased by 36.3%, from 0.1095 rad for HS to 0.0698 rad for SURE, alongside success rates of 40.63% and 95.87%. In G4, both systems completed every attempt, but SURE reduced tilt by 33.1%, from 0.0939 to 0.0628 rad. SURE also had lower action-change rate and torso-angular-velocity CVaR95 in G4 (Appendix Table 9). These results show that its gains in fault traversal were accompanied by reduced large torso deviations.

To test whether the advantage over HS arose solely from switching times, we imposed a shared withdrawal schedule. G3 success remained higher for SURE, at 98.73% versus 31.80% (Appendix B.3). The system-level advantage therefore did not depend solely on selecting different withdrawal times.

Finally, SURE and HS both reached 100% success in G5 after the stream alarm. Both systems could thus complete the task after fallback when an explicit fault signal was available.

E3. The role of physical consistency. We first compare consistency configurations, then separate the vertical module's roles in withdrawal and restoration. Table 2(a) fixes policy weights and the event head. H-only removes vertical withdrawal and all foot-placement verification for restoration. V-only removes horizontal withdrawal, whereas V+collision replaces it with direct collision activation. Both retain the full support-verification procedure before restoring the visual residual. Table 2(b)

progressively restores components of the vertical module to H-only. This separates vertical withdrawal, proprioceptive recontact, and map-based support verification. G5 uses an explicit alarm and is excluded from these ablations.

Table 2: Physical-consistency ablations. H/V denote horizontal/vertical checks. In (a), H-only removes vertical withdrawal and foot-placement verification for restoration. Other rows retain verified restoration. A dash denotes an unevaluated configuration. In (b), groups 1 and 4 match H-only and full SURE. SR and T follow Table 1. Parentheses give evaluation-seed SD in last-digit units (Appendix A.6).

(a) Consistency configurations

Configuration	G1		G2		G3		G4	
	SR ↑	T ↓	SR ↑	T ↓	SR ↑	T ↓	SR ↑	T ↓
H+V (SURE)	**100.00**	0.1069	**100.00**	0.0761	**95.57**	**0.0697**	**100.00**	**0.0628**
H only	**100.00**	**0.1067**	–	–	66.70	0.0756	28.27	0.2439
V only	81.90	0.1652	**100.00**	**0.0760**	10.00	0.0806	**100.00**	**0.0628**
V+collision	**100.00**	0.1074	**100.00**	**0.0760**	87.63	0.0769	**100.00**	0.0629

(b) Vertical module and restoration verification

Group	Withdrawal	Verification	G3		G4	
			SR ↑	T ↓	SR ↑	T ↓
1	H	None	66.70 (87)	0.0756 (2)	28.27 (45)	0.2439 (41)
2	H	Recontact	84.77 (45)	0.0712 (2)	92.13 (47)	0.1220 (14)
3	H	Recontact+map	**95.80** (17)	0.0698 (1)	99.97 (6)	0.0695 (3)
4	**H+V (SURE)**	Recontact+map	95.57 (81)	**0.0697** (1)	**100.00** (0)	**0.0628** (< 1)

Horizontal consistency supplied withdrawal cues that vertical support checks alone missed. With restoration verification unchanged, adding horizontal withdrawal to V-only raised G3 success from 10.00% to 95.57%. Replacing horizontal checks with direct collision activation (V+collision), while retaining vertical checks and the same restoration rules, reduced success from 95.57% to 87.63%. In G1, full SURE likewise reached 100%, compared with 81.90% for V-only. Horizontal map geometry thus improved withdrawal decisions beyond vertical support checks, even when those checks were supplemented by direct collision activation.

Vertical consistency is particularly important under persistent height mismatch. In G4, removing the complete vertical module reduced success from 100% to 28.27% and increased torso tilt from 0.0628 to 0.2439 rad. V-only retained 100% success, showing that support-height checks provide information that horizontal consistency alone cannot replace.

Table 2(b) separates the vertical components responsible for these gains, starting from H-only in group 1. Group 2 adds bilateral proprioceptive recontact, and group 3 additionally verifies support against the map. Map-based support verification raised G4 success from 92.13% to 99.97% and G3 success from 84.77% to 95.80%. Under persistent G4 bias, recontact alone repeatedly permitted visual restoration, whereas map verification kept the residual withdrawn. Recontact confirms physical support but not agreement between that support and the map. Vertical consistency therefore remains necessary during restoration, even when it is not used to trigger withdrawal.

With restoration verification fixed, adding vertical withdrawal improved torso posture under G4. Relative to group 3, SURE reduced torso-tilt CVaR95 from 0.0695 to 0.0628 rad, a 9.6% reduction (Table 2(b)). Velocity-tracking error and torso-angular-velocity CVaR95 also decreased by 7.40% and 27.87%, respectively (Appendix Table 14). In G4, simulator contact records showed no corresponding same-foot lateral contact at the first horizontal withdrawals, whereas SURE first withdrew on vertical contradictions. Vertical consistency thus provides a direct withdrawal cue for height mismatch and constrains unreliable visual re-entry through subsequent support verification.

A recorded G1 encounter illustrates withdrawal followed by support-verified restoration in Appendix Figure 3.

E4. Real-robot validation. Finally, we tested this withdrawal and restoration mechanism on a 29-degree-of-freedom Unitree G1 using measured elevation maps and proprioception. SURE completed box ascent/descent and stair ascent under both nominal and faulty map inputs (Table 3). A representative unmapped-box trajectory showed residual withdrawal after a physical contradiction, followed by restoration after support verification. Appendix A.4 describes the test conditions and robot operation.

Table 3: Real-robot terrain completion under nominal and faulty map input.

Terrain	Map input	Completed / attempted
Box ascent/descent	Nominal	5/5
Box ascent/descent	Box omitted	5/5
Stair ascent	Nominal	5/5
Stair ascent	Delayed by 500 ms	5/5

6 LIMITATIONS

For an unmapped obstacle without a deployable pre-contact cue, SURE's consistency checks cannot prevent the first toe collision because they require feedback from that contact. Between residual withdrawal and restoration, the robot cannot use visual terrain preview to adjust its actions. Continued traversal therefore depends on the blind base policy's terrain capabilities.

Errors in proprioceptive event estimates, foot kinematics, and foot-to-map alignment affect consistency judgments. Shared state-estimation errors can also make incorrect foot and map geometry appear consistent, masking a mismatch. The current foot samples cover a limited range of contact locations and directions and do not capture all lateral or rearward collisions. Calibration, geometric tolerances, and restoration timing may need revalidation across robots or mapping systems.

Consistent local foot interactions do not establish map reliability in regions the robot has not yet encountered. The experiments support performance under the tested terrain and fault conditions, not effectiveness across all terrains.

7 CONCLUSION

SURE compares local map geometry with foot support and collision states inferred from proprioception to control a visual residual over a blind base policy. Observable physical contradictions prompt residual withdrawal, preserving proprioceptive control, while subsequent foot placements must pass support verification before visual corrections resume. Simulations show preserved nominal terrain capability and higher success than the same-weight always-visual policy across four map mismatches without explicit alarms. Horizontal checks improve completion under map delay, while vertical checks improve velocity tracking and torso posture under map-height drift. Map-based support verification further improves completion beyond proprioceptive recontact under delay and height drift, showing that withdrawal and restoration require different physical evidence. On a real humanoid, SURE traverses terrain with nominal and faulty maps and demonstrates consistency-driven withdrawal and restoration. Together, these results establish a physically interpretable route to graceful degradation, retaining terrain adaptation under nominal vision and basic locomotion after perceptual mismatch.

AI USE STATEMENT

AI tools were used to polish the manuscript's wording and LaTeX formatting, and to assist in creating figures and tables.

REFERENCES

Ananye Agarwal, Ashish Kumar, Jitendra Malik, and Deepak Pathak. Legged locomotion in challenging terrains using egocentric vision. In Karen Liu, Dana Kulic, and Jeff Ichnowski (eds.), *Proceedings of The 6th Conference on Robot Learning*, volume 205 of *Proceedings of Machine Learning Research*, pp. 403–415. PMLR, 14–18 Dec 2023. URL https://proceedings.mlr.press/v205/agarwal23a.html.

Woohyun Cha, Junhyeok Cha, Jaeyong Shin, Donghyeon Kim, and Jaeheung Park. Sim-to-real of humanoid locomotion policies via joint torque space perturbation injection. *arXiv preprint arXiv:2504.06585v1*, 2025. URL https://arxiv.org/abs/2504.06585v1.

Péter Fankhauser, Michael Bloesch, Christian Gehring, Marco Hutter, and Roland Siegwart. Robot-centric elevation mapping with uncertainty estimates. In *Mobile Service Robotics*, pp. 433–440. World Scientific, 2014.

Péter Fankhauser, Michael Bloesch, and Marco Hutter. Probabilistic terrain mapping for mobile robots with uncertain localization. *IEEE Robotics and Automation Letters*, 3(4):3019–3026, 2018.

Xinyang Gu, Yen-Jen Wang, and Jianyu Chen. Humanoid-Gym: Reinforcement learning for humanoid robot with zero-shot sim2real transfer. *arXiv preprint arXiv:2404.05695*, 2024a.

Xinyang Gu, Yen-Jen Wang, Xiang Zhu, Chengming Shi, Yanjiang Guo, Yichen Liu, and Jianyu Chen. Advancing humanoid locomotion: Mastering challenging terrains with denoising world model learning. *arXiv preprint arXiv:2408.14472*, 2024b.

Junzhe He, Chong Zhang, Fabian Jenelten, Ruben Grandia, Moritz Bächer, and Marco Hutter. Attention-based map encoding for learning generalized legged locomotion. *Science Robotics*, 10 (105):eadv3604, 2025.

Runhan Huang, Shaoting Zhu, Yilun Du, and Hang Zhao. MoE-Loco: Mixture of experts for multitask locomotion. In *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 14218–14225. IEEE, 2025.

Jemin Hwangbo, Joonho Lee, Alexey Dosovitskiy, Dario Bellicoso, Vassilios Tsounis, Vladlen Koltun, and Marco Hutter. Learning agile and dynamic motor skills for legged robots. *Science robotics*, 4(26):eaau5872, 2019.

Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991.

Tobias Johannink, Shikhar Bahl, Ashvin Nair, Jianlan Luo, Avinash Kumar, Matthias Loskyll, Juan Aparicio Ojea, Eugen Solowjow, and Sergey Levine. Residual reinforcement learning for robot control. In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 6023–6029. IEEE, 2019. doi: 10.1109/ICRA.2019.8794127.

Ashish Kumar, Zipeng Fu, Deepak Pathak, and Jitendra Malik. RMA: Rapid motor adaptation for legged robots. *arXiv preprint arXiv:2107.04034*, 2021.

Junfeng Long, Junli Ren, Moji Shi, Zirui Wang, Tao Huang, Ping Luo, and Jiangmiao Pang. Learning humanoid locomotion with perceptive internal model. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 9997–10003. IEEE, 2025.

Shixin Luo, Songbo Li, Ruiqi Yu, Zhicheng Wang, Jun Wu, and Qiuguo Zhu. PIE: Parkour with implicit-explicit learning framework for legged robots. *IEEE Robotics and Automation Letters*, 9 (11):9986–9993, 2024.

Shihao Ma, Hongjin Chen, Zijun Xu, Yi Zhao, Ke Wu, Ruichen Yang, Leyao Zou, Zhongxue Gan, and Wenchao Ding. CMoE: Contrastive mixture of experts for motion control and terrain adaptation of humanoid robots. *arXiv preprint arXiv:2603.03067*, 2026.

Takahiro Miki, Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning robust perceptive locomotion for quadrupedal robots in the wild. *Science robotics*, 7(62): eabk2822, 2022.

Kasidit Muenprasitivej, Jesse Jiang, Abdulaziz Shamsah, Samuel Coogan, and Ye Zhao. Bipedal safe navigation over uncertain rough terrain: Unifying terrain mapping and locomotion stability. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 11264–11271. IEEE, 2024.

Anderson S. Oliveira, Bryan R. Schlink, W. David Hairston, Peter König, and Daniel P. Ferris. Restricted vision increases sensorimotor cortex involvement in human walking. *Journal of Neurophysiology*, 118(4):1943–1951, 2017. doi: 10.1152/jn.00926.2016.

Ilija Radosavovic, Tete Xiao, Bike Zhang, Trevor Darrell, Jitendra Malik, and Koushil Sreenath. Real-world humanoid locomotion with reinforcement learning. *Science Robotics*, 9(89):eadi9579, 2024.

Junli Ren, Tao Huang, Huayi Wang, Zirui Wang, Qingwei Ben, Junfeng Long, Yanchao Yang, Jiangmiao Pang, and Ping Luo. VB-Com: Learning vision-blind composite humanoid locomotion against deficient perception. *arXiv preprint arXiv:2502.14814*, 2025.

Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*, 2017.

Jonah Siekmann, Kevin Green, John Warila, Alan Fern, and Jonathan Hurst. Blind bipedal stair traversal via sim-to-real reinforcement learning. *arXiv preprint arXiv:2105.08328*, 2021.

Haolin Song, Hongbo Zhu, Tao Yu, Yan Liu, Mingqi Yuan, Wengang Zhou, Hua Chen, and Houqiang Li. Gait-adaptive perceptive humanoid locomotion with real-time under-base terrain reconstruction. *IEEE Robotics and Automation Letters*, 11(4):4969–4976, 2026. doi: 10.1109/LRA.2026.3664167.

Jingkai Sun, Gang Han, Pihai Sun, Wen Zhao, Jiahang Cao, Jiaxu Wang, Qiang Zhang, and Yijie Guo. DPL: Depth-only perceptive humanoid locomotion via realistic depth synthesis and cross-attention terrain reconstruction. *IEEE Robotics and Automation Letters*, 11(8):9407–9414, 2026. doi: 10.1109/LRA.2026.3703584.

Wandong Sun, Baoshi Cao, Long Chen, Yongbo Su, Yang Liu, Zongwu Xie, and Hong Liu. Learning perceptive humanoid locomotion over challenging terrain. In *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 6571–6578. IEEE, 2025.

Dewei Wang, Xinmiao Wang, Xinzhe Liu, Jiyuan Shi, Yingnan Zhao, Chenjia Bai, and Xuelong Li. MoRE: Mixture of residual experts for humanoid lifelike gaits learning on complex terrains. *arXiv preprint arXiv:2506.08840*, 2025a.

Huayi Wang, Zirui Wang, Junli Ren, Qingwei Ben, Tao Huang, Weinan Zhang, and Jiangmiao Pang. BeamDojo: Learning agile humanoid locomotion on sparse footholds. *arXiv preprint arXiv:2502.10363*, 2025b.

Weiji Xie, Chenjia Bai, Jiyuan Shi, Junkai Yang, Yunfei Ge, Weinan Zhang, and Xuelong Li. Humanoid whole-body locomotion on narrow terrain via dynamic balance and reinforcement learning. In *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4751–4758. IEEE, 2025.

Chuanyu Yang, Kai Yuan, Qiuguo Zhu, Wanming Yu, and Zhibin Li. Multi-expert learning of adaptive legged locomotion. *Science Robotics*, 5(49):eabb2174, 2020.

Ziwon Yoon, Lawrence Y Zhu, Jingxi Lu, Lu Gan, and Ye Zhao. STATE-NAV: Stability-aware traversability estimation for bipedal navigation on rough terrain. *IEEE Robotics and Automation Letters*, 11(2):2338–2345, 2026. doi: 10.1109/LRA.2025.3648502.

Qiang Zhang, Gang Han, Jingkai Sun, Wen Zhao, Chenghao Sun, Jiahang Cao, Jiaxu Wang, Yijie Guo, and Renjing Xu. Distillation-PPO: A novel two-stage reinforcement learning framework for humanoid robot perceptive locomotion. In *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2916–2922. IEEE, 2025a.

Zhikai Zhang, Jun Guo, Chao Chen, Jilong Wang, Chenghuai Lin, Yunrui Lian, Han Xue, Zhenrong Wang, Maoqi Liu, Jiangran Lyu, et al. Track any motions under any disturbances. *arXiv preprint arXiv:2509.13833*, 2025b.

Shaoting Zhu, Ziwen Zhuang, Mengjie Zhao, Kun-Ying Lee, and Hang Zhao. Hiking in the wild: A scalable perceptive parkour framework for humanoids. *arXiv preprint arXiv:2601.07718*, 2026.

Ziwen Zhuang, Zipeng Fu, Jianren Wang, Christopher Atkeson, Soeren Schwertfeger, Chelsea Finn, and Hang Zhao. Robot parkour learning. *arXiv preprint arXiv:2309.05665*, 2023.

Ziwen Zhuang, Shenzhe Yao, and Hang Zhao. Humanoid parkour learning. *arXiv preprint arXiv:2406.10759*, 2024.

A METHOD AND EXPERIMENTAL SETUP

A.1 DEPLOYMENT INFORMATION BOUNDARY

Table 4: Deployment inputs and training supervision.

Signal	Action policy	Consistency	Training labels
Commands and proprioceptive history	yes	yes	no
Received elevation map	yes	yes	no
Current-frame kinematic foot geometry	no	yes	no
Query bounds and numerical validity	no	yes	no
Simulated terrain and contact/collision labels	no	no	yes
Explicit G5 flag	SURE and HS latch blind fallback		no

The proprioceptive event head receives only proprioceptive history and derived kinematics. Geometry uses current joint angles, gravity direction, and the received map. Simulator ground truth supports training and evaluation but does not enter the deployed controller.

A.2 ACTION NETWORKS AND TRAINING

The action networks combine multiple experts through a learned gate (Jacobs et al., 1991; Shazeer et al., 2017), following mixture-of-experts approaches to multi-skill locomotion (Yang et al., 2020; Huang et al., 2025; Wang et al., 2025a; Ma et al., 2026). The proprioceptive encoder flattens five history frames and applies a multilayer perceptron (MLP) with widths $512 \to 256 \to 128$. The map encoder processes a 32×40 elevation map using three convolutional layers, with channels 16/32/32, kernel sizes 5/3/3, and strides 2/2/1. An MLP then projects the output to 128 visual features. Each action expert has hidden widths of 256 and 128 and outputs 29 action dimensions. A router with one 128-unit hidden layer forms a weighted combination of expert outputs.

Visual-availability curriculum and gradients. Policy training lasts 30,000 iterations. Each episode samples a visual-availability mask m_e, fixed until its end, yielding the action $\mathbf{a}_{\text{blind}} + m_e \Delta \mathbf{a}_{\text{visual}}$. At iteration k, the blind-only probability is

$$p_{\text{blind}}(k) = \Pr(m_e = 0) = \begin{cases} 1, & k < 15000, \\ 1 - 0.9(k - 15000)/5000, & 15000 \le k < 20000, \\ 0.1, & k \ge 20000. \end{cases} \tag{6}$$

The first half is blind-only, followed by a linear increase in visual availability. The final blind-to-vision episode sampling ratio is $1 : 9$.

The visual branch applies stop-gradient to its proprioceptive-feature and blind-action inputs. The blind branch still receives PPO gradients through the combined action. At deployment, physical consistency controls β. The proprioceptive event head is trained and calibrated separately (Appendix C.1).

Optimization. All four policies optimize Equation (1) with PPO and an asymmetric actor–critic. The actor uses deployable observations, while the critic additionally receives privileged observations. The critic MLP has hidden widths of 512, 256, and 128, with LayerNorm and ELU in each layer. Each iteration collects 24 steps per environment and uses five learning epochs, eight mini-batches, and Adam. The initial learning rate is 10^{-3} and adapts to Kullback–Leibler (KL) divergence. The PPO clip coefficient is 0.2, discount factor 0.97, and generalized advantage estimation coefficient 0.95. The entropy coefficient and target KL divergence are both 0.01. The maximum gradient norm is 1 and the router load-balancing coefficient is 0.01.

Map randomization. SURE and Visual add uniform height noise $\mathcal{U}(-0.03, 0.03)$ to the map. Heavy-DR additionally replaces a local block with probability 0.15 per episode. Block width and height each span 15–25 cells, with replacement values sampled from $\mathcal{U}(-1, 1)$. Heavy-DR also introduces random observation delays of 0.1–0.5 s. SURE training does not use the G1–G5 fault-injection protocols. The combined action is clipped to $\mathbf{a}_t^{\text{cmd}}$ and converted to joint targets $\mathbf{q}_t^{\text{des}} = \mathbf{q}^{\text{ref}} + \boldsymbol{\alpha} \odot \mathbf{a}_t^{\text{cmd}}$. A low-level PD controller tracks these targets.

Table 5: Action-policy training configurations. Expert counts are listed as blind/visual.

Policy	Experts	Visual availability	Map perturbation and delay
SURE	4 / 4	Episode curriculum	Uniform noise, no added delay
Visual	0 / 8	Always enabled	Uniform noise, no added delay
Heavy-DR	0 / 8	Always enabled	Uniform noise and block replacement, 0.1–0.5 s delay
Blind	8 / 0	No vision	Map-independent

A.3 Geometric tests and control state

Coordinates and foot samples. Current joint angles determine foot poses, and samples are transformed into the gravity-aligned local map frame centered at the base. The elevation map has 32×40 cells at 0.04 m resolution. Geometry queries read the cell containing each sample, after removing the 0.76 m offset used in map encoding.

Each foot uses five sole points for support-height checks and three front-surface points for horizontal geometry. The samples are fixed relative to the foot and transformed with its pose. Table 6 gives their local coordinates.

Table 6: Foot samples in the local ankle frame (m), shared by both feet.

Sole points			Front points		
x	y	z	x	y	z
0.132	0	−0.035	0.108	−0.032	−0.025
−0.054	0	−0.035	0.108	0.032	−0.025
0.039	−0.018	−0.035	0.142	0	−0.025
0.039	0.018	−0.035			
0.039	0	−0.035			

Consistency tests. The per-foot tests follow Equations (3)–(5). The support tolerance $\epsilon_s \approx 0.062$ m combines map-training noise with an allowance for discrete sole sampling. The horizontal geometry margin is $\epsilon_g = 0.03$ m. Matching one sole sample to map height is sufficient to accommodate partial support. All conditions share these parameters, and out-of-bounds or invalid queries produce no contradiction.

Withdrawal and restoration. We add time subscript t to the per-foot variables and define

$$c_t^v = \max_f c_{f,t}^v, \qquad c_t^h = \max_f c_{f,t}^h. \tag{7}$$

The control interval is $\Delta t = 0.02$ s, and horizontal contradictions are held for two control steps. Either contradiction can request withdrawal of the visual residual. A new withdrawal request is deferred when neither support nor collision feedback is present.

After withdrawal, and while observations remain valid, each foot must pass through non-support and subsequently establish support consistent with the map. Once contradictions clear and both feet pass verification, continuous confirmation for 0.2 s is required before setting $\beta_t = 1$. The explicit G5 stream-fault flag instead keeps SURE and HS in blind control until episode reset.

A.4 Evaluation configurations

SURE and its fixed $\beta = 0/1$ endpoints share action-policy weights. Blind, Visual, and Heavy-DR Visual use the training configurations in Table 5. HS switches between the independently trained blind and visual policies,

$$\mathbf{a}_t^{\mathrm{HS}} = (1 - b_t)\pi_{\mathrm{B}}(\mathbf{o}_t^P) + b_t \pi_{\mathrm{V}}(\mathbf{o}_t^P, \widetilde{H}_t), \qquad b_t \in \{0, 1\}. \tag{8}$$

Its proprioceptive event head is matched to these policies. HS uses the same family of geometric checks, withdrawal rules, and restoration rules as SURE. The fixed endpoints, Visual, and Heavy-DR Visual do not use consistency gating.

Real-robot tests. The robot is a 29-degree-of-freedom Unitree G1 equipped with a Livox MID-360 LiDAR. FAST-LIO estimates pose, and Elevation Mapping CuPy constructs the elevation map. At each startup, the system rescans the starting area and initializes a new map. The near-body sensing blind region is filled with flat-ground height. The policy controls the robot in closed loop using elevation maps and proprioception.

Tests comprise ascent/descent of a 16 cm-high box and ascent of a three-step staircase with 16 cm risers and 33 cm tread depth. Box tests use nominal maps or maps omitting the box. Stair tests use nominal maps or maps delayed by 500 ms. Physical terrain remains unchanged, and proprioceptive observations update in real time. Each condition has five trials, with success defined as completing box ascent/descent or reaching the stair top (Table 3).

A representative unmapped-box trajectory shows residual withdrawal followed by restoration after subsequent foot placements pass support verification. In nominal box trials, map noise near the vertical face at the far edge also caused brief withdrawal. The supplementary video shows real-robot walking and the residual withdrawal and restoration process.

A.5 Training terrains and test conditions

Nominal testing covers nine conditions, including rough ground, randomly placed low obstacles, gutters, and ascent and descent on slopes, high steps, and stairs. Table 7 specifies their geometry. Training also includes a thin-tread stair variant, with tread thickness and overhang both 0.02 m, grouped under stairs.

Table 7: Training ranges and nominal test parameters. Units are meters except for angles and counts. Randomized terrains list their fixed sampling ranges.

Terrain	Training range	Test parameters
T0 Rough ground	Height amplitude 0.005–0.025	Height $[-0.015, 0.015]$, increment 0.005
T1 Slopes	Incline 5°–25°	Incline 15°, ascent/descent
T2 High steps	Height 0.20–0.50, tread width 1.00	Height 0.35, tread width 1.00, ascent/descent
T3 Stairs	Height 0.10–0.25, tread depth 0.32–0.40 (thin-tread variant 0.30–0.40)	Height 0.175, tread depth 0.36, 16 steps, ascent/descent
T4 Low obstacles	Height 0.05–0.40, length 0.40–1.264, width 0.30–0.948, count 50–140, minimum spacing 0.05–0.25	Height 0.05–0.20, length 0.40–1.048, width 0.30–0.732, count 80, minimum spacing 0.20
T5 Repeated gutters	Width 0.20–0.50, depth 4.00, flat interval 1.00	Width 0.35, depth 4.00, flat interval 1.00

Nominal terrains span 15×15 m with a 2 m-wide central starting platform. The forward command is 0.75 m/s, with heading control capped at 0.3 rad/s. Physics and control steps are 0.005 s and 0.02 s, respectively.

G1 uses a protrusion of length, width, and height $0.30 \times 0.50 \times 0.20$ m, centered 1.25 m ahead, while the map depicts flat ground. G2 uses a depression of length, width, and depth $1.00 \times 0.50 \times 0.20$ m, centered 1.00 m ahead, while the map depicts a continuous plane. G3 uses T3 ascent with equally represented delays of 0.1–1.0 s in 0.1 s increments, each fixed within an episode. G4 uses 20×20 m flat ground. Map height remains unchanged for 2 s, decreases by 0.30 m over 2–4 s, and then retains this bias while the physical terrain remains unchanged. G5 uses the same ascending stairs and stops map updates after 1.20 m of forward progress. After 0.60 s, it zeroes the map encoding and reports a stream-fault flag, causing SURE and HS to fall back to blind control.

A.6 METRICS AND STATISTICS

Task success. An episode succeeds if the robot reaches the 5 m finish within 20 s without earlier termination from falling or contact outside the feet and ankles. Success rate is the fraction of attempted episodes that succeed.

Motion metrics. Let $\mathcal{I}_e$ contain episode e's valid control steps, with $n_e = |\mathcal{I}_e|$. Velocity-tracking error E_v measures planar command following,

$$E_{v,e} = \frac{1}{n_e} \sum_{t \in \mathcal{I}_e} \|\mathbf{v}_{xy,t}^{\mathrm{cmd}} - \mathbf{v}_{xy,t}^{\mathrm{base}}\|_2. \tag{9}$$

Torso tilt T measures deviation from upright, while torso angular velocity W measures rotational speed. Both use CVaR95, the mean of the largest 5% of values within an episode,

$$T_e = \mathrm{CVaR}_{95}\left[\arccos\left(\frac{-g_{z,t}^{\mathrm{torso}}}{\|\mathbf{g}_t^{\mathrm{torso}}\|_2}\right)\right], \qquad W_e = \mathrm{CVaR}_{95}\left[\|\boldsymbol{\omega}_t^{\mathrm{torso}}\|_2\right]. \tag{10}$$

Here $\mathbf{g}^{\mathrm{torso}}$ is gravity expressed in the torso frame, and $\boldsymbol{\omega}^{\mathrm{torso}}$ is torso angular velocity. Action-change rate D_a measures changes in executed actions between adjacent control steps,

$$D_{a,e} = \sqrt{\frac{1}{D|\mathcal{J}_e|} \sum_{t \in \mathcal{J}_e} \left\| \frac{\mathbf{a}_t^{\mathrm{cmd}} - \mathbf{a}_{t-1}^{\mathrm{cmd}}}{\Delta t} \right\|_2^2}, \qquad \mathcal{J}_e = \{t \in \mathcal{I}_e : t-1 \in \mathcal{I}_e\}. \tag{11}$$

The executed normalized action is $\mathbf{a}^{\mathrm{cmd}}$, with dimension D. Units for E_v, T, W, and D_a are m/s, rad, rad/s, and s^{-1}, respectively. Visual-on fraction is the episode-level time average of β_t for SURE or b_t for HS.

Aggregation. Simulation control evaluations use three evaluation seeds per configuration and terrain or fault condition, with 1,000 episodes per seed and fixed policy weights. Clean weights the nine nominal conditions equally. Motion metrics are computed per episode and averaged over valid episodes, including valid trajectories up to failure. The notation $\pm$ denotes sample standard deviation (SD) across evaluation seeds. Compact tables show SD in last-digit units, for example $99.97(6) = 99.97 \pm 0.06$. A positive SD smaller than one last-digit unit is shown as < 1. Bold marks the best displayed mean per condition and metric, including ties.

B SUPPLEMENTARY COMPARISONS

B.1 TERRAIN-LEVEL SUCCESS AND MOTION METRICS

Table 8 reports terrain-level success. Table 9 summarizes success and motion metrics for Clean and G1–G5. Definitions and aggregation follow Appendix A.6.

Table 8: Nominal-terrain success (%, ↑). The last two columns are SURE's fixed endpoints.

Terrain	SURE	Blind	Visual	Heavy-DR	HS	$\beta = 0$	$\beta = 1$
T0	**100.00** (0)	**100.00** (0)	**100.00** (0)	**100.00** (0)	**100.00** (0)	**100.00** (0)	**100.00** (0)
T1 up	**100.00** (0)	**100.00** (0)	**100.00** (0)	**100.00** (0)	**100.00** (0)	**100.00** (0)	**100.00** (0)
T1 down	**100.00** (0)	**100.00** (0)	**100.00** (0)	**100.00** (0)	**100.00** (0)	**100.00** (0)	**100.00** (0)
T2 up	**100.00** (0)	0.00 (0)	**100.00** (0)	0.00 (0)	96.13 (70)	0.00 (0)	**100.00** (0)
T2 down	**100.00** (0)	**100.00** (0)	**100.00** (0)	97.30 (26)	**100.00** (0)	**100.00** (0)	**100.00** (0)
T3 up	**100.00** (0)	**100.00** (0)	**100.00** (0)	**100.00** (0)	99.97 (6)	**100.00** (0)	**100.00** (0)
T3 down	**100.00** (0)	**100.00** (0)	**100.00** (0)	**100.00** (0)	**100.00** (0)	**100.00** (0)	**100.00** (0)
T4	**100.00** (0)	20.23 (131)	99.60 (0)	95.93 (45)	89.13 (117)	**100.00** (0)	**100.00** (0)
T5	99.33 (15)	0.83 (6)	**100.00** (0)	97.20 (50)	94.27 (90)	0.37 (6)	99.93 (6)

Table 9: Success and motion metrics for the seven configurations.

Condition	Configuration	SR (%) ↑	$T \downarrow$	$E_v \downarrow$	$D_a \downarrow$	$W \downarrow$
Clean	SURE	99.93	0.0733	0.0877	14.913	1.763
Clean	Blind	69.01	**0.0579**	0.1468	14.499	1.316
Clean	Visual	99.96	0.0854	**0.0669**	15.074	**1.259**
Clean	Heavy-DR Visual	87.83	0.0927	0.1596	16.676	1.938
Clean	HS	97.72	0.0854	0.0712	15.186	1.303
Clean	SURE $\beta = 0$	77.82	0.0789	0.1352	**13.995**	1.676
Clean	SURE $\beta = 1$	**99.99**	0.0753	0.0862	14.584	1.687
G1	SURE	**100.00**	0.1071	0.0705	15.989	2.566
G1	Blind	**100.00**	0.0887	0.0796	15.926	1.628
G1	Visual	0.00	0.0811	0.0846	17.043	**1.535**
G1	Heavy-DR Visual	0.00	**0.0794**	0.0913	17.470	1.606
G1	HS	99.67	0.1115	**0.0615**	16.928	2.287
G1	SURE $\beta = 0$	**100.00**	0.1119	0.1016	**15.342**	2.668
G1	SURE $\beta = 1$	79.27	0.1548	0.0815	15.813	2.973
G2	SURE	**100.00**	0.0761	**0.0664**	16.062	2.136
G2	Blind	**100.00**	0.0894	0.0755	15.604	**1.560**
G2	Visual	0.00	0.1013	0.2055	21.707	3.659
G2	Heavy-DR Visual	7.33	0.0955	0.1280	18.972	2.049
G2	HS	99.90	0.0878	0.0744	17.254	1.928
G2	SURE $\beta = 0$	**100.00**	**0.0639**	0.0899	**15.325**	2.195
G2	SURE $\beta = 1$	28.67	0.1339	0.3702	20.209	3.030
G3	SURE	95.87	0.0698	0.1054	20.313	2.488
G3	Blind	**100.00**	0.0675	0.0963	18.345	**1.426**
G3	Visual	10.03	0.1221	0.2864	15.854	2.541
G3	Heavy-DR Visual	96.33	0.0705	**0.0936**	17.995	1.867
G3	HS	40.63	0.1095	0.2242	20.986	2.462
G3	SURE $\beta = 0$	**100.00**	**0.0592**	0.1041	19.500	2.474
G3	SURE $\beta = 1$	10.00	0.0803	0.2033	**15.531**	2.652
G4	SURE	**100.00**	0.0628	0.1000	12.412	1.062
G4	Blind	**100.00**	**0.0449**	0.1048	14.028	1.021
G4	Visual	0.00	0.3007	0.4968	23.135	4.956
G4	Heavy-DR Visual	**100.00**	0.0642	0.3661	17.497	1.370
G4	HS	**100.00**	0.0939	**0.0842**	14.806	1.407
G4	SURE $\beta = 0$	**100.00**	0.0565	0.1041	**12.119**	**0.954**
G4	SURE $\beta = 1$	0.10	0.3853	0.6106	28.926	5.533
G5	SURE	**100.00**	0.0618	0.0989	19.513	2.397
G5	Blind	**100.00**	0.0676	0.0963	**18.344**	**1.426**
G5	Visual	0.00	0.2729	0.3220	21.314	4.418
G5	Heavy-DR Visual	0.00	0.0754	0.1970	18.622	1.855
G5	HS	**100.00**	0.0771	**0.0891**	18.636	1.578
G5	SURE $\beta = 0$	**100.00**	**0.0592**	0.1040	19.503	2.476
G5	SURE $\beta = 1$	0.00	0.1026	0.1697	18.635	2.529

B.2 VISUAL-ON FRACTION

Table 10: Visual-on time fraction (%). SURE enables the visual residual, whereas HS selects the complete visual policy.

Configuration	Clean	G1	G2	G3	G4	G5
SURE	96.16	82.66	72.88	33.14	18.46	26.23
HS	94.81	84.76	76.64	59.62	19.46	26.85

B.3 SHARED WITHDRAWAL SCHEDULE

To remove differences in withdrawal timing, this comparison bypasses consistency triggers and imposes the same visual-control schedule on SURE and HS. Both withdraw vision 1.34 s after forward commands begin. Vision resumes at 3.34 s on nominal stairs but remains off in G3. Table 11 also reports always-blind and always-visual controls.

Table 11: Shared withdrawal schedule and fixed visual states. Clean denotes nominal stair ascent. Bold compares SURE and HS under the same condition and visual-control schedule.

Condition	Visual control	System	SR (%) ↑	$T\downarrow$	$E_v\downarrow$	$D_a\downarrow$	$W\downarrow$
Clean	Always blind	SURE	**100.00**	**0.0592**	0.1041	19.502	2.475
Clean	Always blind	HS	**100.00**	0.0675	**0.0963**	**18.349**	**1.425**
Clean	Always visual	SURE	**100.00**	**0.0565**	0.0903	18.268	2.356
Clean	Always visual	HS	**100.00**	0.0739	**0.0659**	**17.008**	**1.294**
Clean	Shared schedule	SURE	**100.00**	**0.0670**	0.0957	**19.541**	2.428
Clean	Shared schedule	HS	**100.00**	0.1016	**0.0843**	19.712	**2.164**
G3	Always blind	SURE	**100.00**	**0.0592**	0.1040	19.503	2.476
G3	Always blind	HS	**100.00**	0.0675	**0.0963**	**18.349**	**1.425**
G3	Always visual	SURE	10.00	**0.0804**	**0.2032**	**15.561**	2.657
G3	Always visual	HS	**10.03**	0.1217	0.2863	15.852	**2.518**
G3	Shared schedule	SURE	**98.73**	**0.0659**	**0.1012**	**19.886**	**2.462**
G3	Shared schedule	HS	31.80	0.1192	0.2587	20.728	2.780

C EVENT HEADS AND CONSISTENCY ABLATIONS

C.1 PROPRIOCEPTIVE EVENT HEADS

The event head receives five proprioceptive history frames, each containing 57 features. Inputs include 12 FK-derived sole coordinates and three components each of body angular velocity, projected gravity, and motion commands. Positions, velocities, and previous actions of the 12 leg joints supply the remaining 36 features. A gated recurrent unit (GRU) with 64 hidden units outputs support and sustained-collision probabilities for each foot. Simulated contact and collision labels provide supervision only, and all network inputs are available at deployment.

SURE and HS each collect rollouts from their blind and visual endpoints without injected visual faults. Complete episodes are split into Train, Val, and Test in an approximately $7:2:1$ ratio. Val contains disjoint model-selection and calibration subsets. Each head is trained for 200 epochs using Adam, a 10^{-3} learning rate, and batches of 65,536 frames. The loss equally weights support and collision binary cross-entropies, with additional weighting of positive collision labels.

After training, temperature scaling and event thresholds are determined on the calibration subset. For foot f, support and collision logits ℓ_f^{sup} and ℓ_f^{col} give calibrated probabilities

$$p_f^{\mathrm{sup}} = \sigma\Big(\ell_f^{\mathrm{sup}}/T_{\mathrm{sup}}\Big), \qquad p_f^{\mathrm{col}} = \sigma\big(\ell_f^{\mathrm{col}}/T_{\mathrm{col}}\big), \tag{12}$$

where $T_{\mathrm{sup}}, T_{\mathrm{col}} > 0$ are temperature parameters. Thresholding gives the main-text states s_f and e_f. Network and calibration parameters then remain fixed. Table 12 reports predictions on Test.

C.2 CONSISTENCY ABLATIONS

All variants share the action policy and proprioceptive event head. H+V uses both checks with verified restoration. H removes vertical withdrawal and all foot-placement verification for restoration, retaining horizontal checks and common temporal rules. V removes horizontal withdrawal, while V+collision replaces it with direct proprioceptive collision activation. Both retain bilateral recontact and map-based support verification before restoring the residual. The complete-module H configuration was not evaluated in G2.

Table 13 gives success and motion metrics for Clean and G1–G4. Clean compares H+V and V+collision. Appendix C.3 separately compares four restoration-verification conditions.

Table 12: Event-prediction precision P and recall R (%, ↑). L/R denote left/right feet. HS endpoints use the independently trained policies.

Endpoint	Foot	Support		Collision	
		P	R	P	R
SURE $\beta = 0$	L	99.59	99.35	96.75	72.33
	R	99.57	99.51	89.72	74.62
SURE $\beta = 1$	L	99.78	99.78	89.71	61.14
	R	99.80	99.77	67.16	74.70
HS blind	L	99.78	99.63	97.13	75.61
	R	99.61	99.65	89.89	72.81
HS visual	L	99.88	99.82	58.96	71.08
	R	99.84	99.80	80.63	64.86

Table 13: Consistency ablations. H excludes vertical withdrawal and foot-placement verification for restoration. Clean aggregates nine terrains, and G4 denotes relative map-height drift.

Condition	Configuration	SR (%) ↑	$T \downarrow$	$E_v \downarrow$	$D_a \downarrow$	$W \downarrow$
Clean	H+V	**99.95**	**0.0732**	**0.0877**	14.912	1.763
Clean	V+collision	99.89	**0.0732**	**0.0877**	**14.874**	**1.758**
G1	H+V	**100.00**	0.1069	**0.0704**	15.994	2.562
G1	H	**100.00**	**0.1067**	0.0715	16.409	**2.532**
G1	V	81.90	0.1652	0.0881	16.554	3.284
G1	V+collision	**100.00**	0.1074	0.0705	**15.979**	2.568
G2	H+V	**100.00**	0.0761	**0.0664**	16.062	2.136
G2	V	**100.00**	**0.0760**	**0.0664**	**16.059**	**2.135**
G2	V+collision	**100.00**	**0.0760**	**0.0664**	**16.059**	**2.135**
G3	H+V	**95.57**	**0.0697**	**0.1058**	20.263	**2.488**
G3	H	66.70	0.0756	0.1412	21.100	2.642
G3	V	10.00	0.0806	0.2031	**15.779**	2.657
G3	V+collision	87.63	0.0769	0.1270	20.343	2.555
G4	H+V	**100.00**	**0.0628**	**0.1000**	12.415	1.062
G4	H	28.27	0.2439	0.3819	25.499	4.677
G4	V	**100.00**	**0.0628**	**0.1000**	12.417	**1.061**
G4	V+collision	**100.00**	0.0629	**0.1000**	**12.404**	1.063

C.3 VERTICAL-MODULE AND RESTORATION ABLATIONS

This comparison varies withdrawal cues and restoration verification separately. Group numbers match Table 2(b). All groups share the action policy and proprioceptive event head, retaining withdrawal holding and continuous 0.2 s restoration confirmation.

1. **Horizontal withdrawal without foot-placement verification.** Horizontal contradictions trigger withdrawal. After they clear, common temporal rules permit restoration without bilateral recontact or map-height checks.
2. **Horizontal withdrawal with proprioceptive recontact.** Withdrawal follows group 1. Restoration additionally requires proprioceptive confirmation that each foot has transitioned from non-support to support, without checking map height.
3. **Horizontal withdrawal with recontact and map verification.** Withdrawal is unchanged. In addition to group 2's requirements, each new foot placement must pass a sole-to-map height-consistency check.
4. **Full SURE.** Restoration follows group 3, while either horizontal or vertical contradictions trigger withdrawal.

Conditions comprise flat Clean, G1 unmapped protrusion, G3 stale stair maps, and G4 relative map-height drift. Flat Clean is G4's no-drift reference, distinct from the nine-terrain main-text aggregate. The four groups are compared within one batch. Groups 1 and 4 also supply the matched H and H+V

results for G1, G3, and G4 in Appendix C.2. All groups achieve 100% success in flat Clean and G1. G3/G4 success appears in Table 2(b), and Table 14 gives supplementary motion metrics.

Table 14: Motion metrics for vertical-module and restoration ablations. Groups 1–4 follow this subsection and Table 2(b).

Condition	Group	$T \downarrow$	$E_v \downarrow$	$D_a \downarrow$	$W \downarrow$
Flat Clean	1	**0.0597** (< 1)	**0.0685** (< 1)	12.395 (7)	**1.236** (< 1)
Flat Clean	2	0.0598 (1)	**0.0685** (< 1)	**12.392** (8)	**1.236** (< 1)
Flat Clean	3	0.0598 (< 1)	**0.0685** (< 1)	12.395 (3)	**1.236** (< 1)
Flat Clean	4	0.0598 (< 1)	**0.0685** (< 1)	12.396 (10)	**1.236** (< 1)
G1	1	0.1067 (8)	0.0715 (5)	16.409 (36)	2.532 (20)
G1	2	**0.0905** (5)	**0.0626** (4)	**15.582** (22)	**2.227** (6)
G1	3	0.1075 (2)	0.0704 (1)	15.993 (6)	2.564 (5)
G1	4	0.1069 (4)	0.0704 (1)	15.994 (11)	2.562 (5)
G3	1	0.0756 (2)	0.1412 (9)	21.100 (52)	2.642 (7)
G3	2	0.0712 (2)	0.1136 (2)	20.283 (18)	**2.478** (1)
G3	3	0.0698 (1)	**0.1055** (2)	20.315 (12)	2.490 (2)
G3	4	**0.0697** (1)	0.1058 (8)	**20.263** (40)	2.488 (2)
G4	1	0.2439 (41)	0.3819 (67)	25.499 (132)	4.677 (55)
G4	2	0.1220 (14)	0.1364 (15)	17.870 (65)	2.631 (22)
G4	3	0.0695 (3)	0.1080 (7)	13.229 (32)	1.472 (8)
G4	4	**0.0628** (< 1)	**0.1000** (< 1)	**12.415** (6)	**1.062** (1)

Figure 3 illustrates residual withdrawal and subsequent restoration during a G1 encounter.

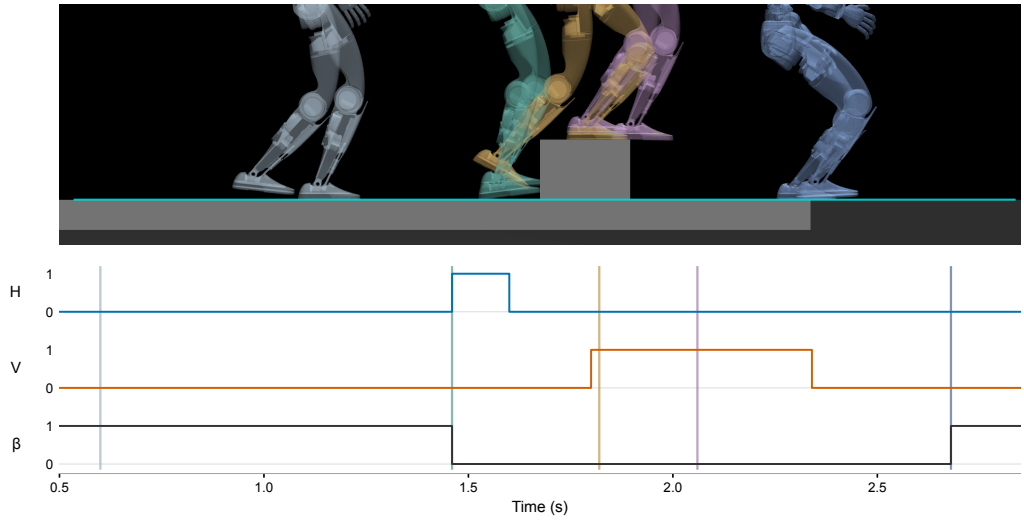


Figure 3: Withdrawal and restoration during a G1 encounter. The upper panel overlays foot poses, mapped heights (cyan), and terrain (gray). The lower panel shows horizontal contradictions H, vertical contradictions V, and residual switch β. Colors link poses to time guides.