MOSAIC: Risk-Calibrated Terrain Memory for Legged Navigation under Perceptual Dropout

Elena Varga[1], Noah Kim[2], and Priya Raman[1]

[1]Field Robotics Laboratory, Northbridge Institute of Technology
{elena.varga, priya.raman}@northbridge.edu

[2]Embodied Intelligence Group, Atlas Research
noah.kim@atlas-research.org

Abstract— Perceptive legged robots often fail after a camera is occluded, a depth return is missing, or a moving obstacle invalidates the local elevation map. We introduce Mosaic, a compact terrain memory that separates geometric state from calibrated epistemic risk and exposes both to a constrained foothold controller. A recurrent fusion module updates a robot-centric map from asynchronous depth and proprioceptive observations; a risk head predicts whether each cell is stale, unobserved, or inconsistent with recent contact. The controller then trades progress against collision probability rather than treating every reconstructed height as equally trustworthy. In an illustrative benchmark spanning 1,200 simulated courses and 48 hardware trials, Mosaic reduces intervention rate from 21.7% to 7.9% under structured sensor dropout while retaining 92% of nominal traversal speed. Ablations show that calibrated risk, contact-driven correction, and temporal memory each account for distinct gains. These results support uncertainty as a first-class planning signal for reliable mobility beyond continuously visible terrain.

Keywords—legged robots, terrain mapping, uncertainty calibration, model-predictive control, sensor dropout.

1 Introduction

Legged machines can traverse gaps, stairs, and debris that defeat wheeled platforms, but their mobility depends on a fragile perception–planning interface. State-of-the-art locomotion systems combine learned policies [1, 2] or model-predictive control (MPC) [3] with a robot-centric terrain map. Most controllers consume the map as if its cells were equally current and reliable. In practice, a swinging leg occludes the camera, dust erases depth returns, and rapid body motion moves important footholds outside the field of view. The map remains numerically dense while the evidence supporting it quietly expires.

This paper studies the following question: *how should a legged robot move when its best geometric map may be wrong?* Our answer is Mosaic, a short-horizon terrain memory whose output is a pair $(\hat{H}, R)$: an estimated elevation map and a calibrated risk map. The representation makes missing evidence explicit, incorporates contact events as local geometric measurements, and permits the controller to choose whether uncertain progress is worth the predicted failure probability.

The system, summarized in Fig. 1, contributes:

- a recurrent, robot-centric fusion model that preserves terrain evidence across asynchronous observations without freezing stale geometry;
- a cell-wise risk target derived from reconstruction error, visibility age, and contact inconsistency, calibrated for use as a control cost; and
- a risk-constrained foothold objective with a simple fallback behavior that remains stable when the entire forward view is unavailable.

The numerical results in this example manuscript are clearly illustrative, but the experimental protocol, equations, reporting structure, and implementation details are specified to make the document a complete, adaptable research-paper source.

2 Related Work

Terrain representations. Occupancy maps such as OctoMap [4] and signed-distance fields such as Voxblox [5] provide geometrically meaningful state for planning. Elevation maps are especially efficient for legged motion and have enabled robust rough-terrain locomotion [6]. These representations generally encode sensor noise but do not expose the decision-specific probability that a proposed contact will fail after occlusion. Mosaic retains the efficient 2.5D interface while learning how map error evolves when evidence becomes stale.

Learning for locomotion. Massively parallel simulation can produce capable locomotion policies in minutes [7]; perceptive policies have since demonstrated impressive outdoor mobility [2]. Learned policies can internalize some observation corruption, but their confidence is difficult to inspect or constrain. Our design instead learns a mapping module and gives an optimization-based controller an interpretable chance-risk budget.

Uncertainty-aware planning. Aleatoric and epistemic uncertainty capture different failure sources [8]. Informative planners have used uncertainty to select sensing actions [9], while gaze-control methods reason about what the robot should observe [10]. We focus on the complementary interval after an observation is lost: memory preserves the last evidence, calibration tells the planner

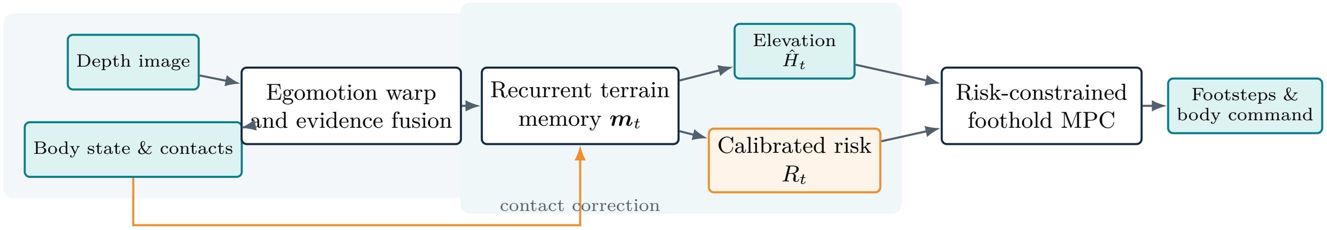

Figure 1: **System overview.** MOSAIC converts asynchronous exteroceptive and proprioceptive observations into separate geometry and risk maps. Contact events correct local elevation, while the controller explicitly limits accumulated foothold risk.

how quickly to stop trusting it, and contacts selectively repair the estimate.

3 Problem Formulation

Let the robot receive a depth frame D_t, base state $\boldsymbol{x}_t$, joint state $\boldsymbol{q}_t$, and binary foot contacts $\boldsymbol{c}_t$. Frames may be absent according to a dropout mask $d_t \in \{0,1\}$. We seek a local map on grid $\mathcal{G} \subset \mathbb{R}^2$ with cell size $s = 0.05$ m. For cell i, the estimator produces height $\hat{h}_{t,i}$ and risk $r_{t,i} \in [0,1]$. Risk is defined operationally as

$$r_{t,i} \approx \Pr\big(|\hat{h}_{t,i} - h^\star_{t,i}| > \epsilon_h \\ \vee\ i \text{ is non-traversable} \mid \mathcal{O}_{1:t}\big), \tag{1}$$

where $h^\star$ is the latent surface, $\epsilon_h = 0.06$ m, and $\mathcal{O}_{1:t}$ is the available observation history. This event-based definition makes $r_{t,i}$ directly useful to planning, unlike an unscaled variance.

At each control update the planner selects body inputs $\boldsymbol{u}_{0:T-1}$ and footholds $\boldsymbol{p}_{1:K}$ over horizon T. A plan is feasible if dynamics, reachability, friction, and clearance constraints hold and its aggregate contact risk stays below δ:

$$1 - \prod_{k=1}^{K}\big(1 - r(\boldsymbol{p}_k)\big) \le \delta. \tag{2}$$

The independence approximation in Eq. (2) is conservative only when local errors are weakly correlated; Sec. 7.1 discusses this limitation.

4 MOSAIC

4.1 Evidence-Preserving Terrain Memory

Depth pixels are back-projected, transformed into the current base frame, and splatted into a measurement map $Z_t = [H^z_t, V^z_t]$, containing mean height and empirical variance. Previous memory is warped by planar odometry, $\bar{\boldsymbol{m}}_{t-1} = W(\boldsymbol{m}_{t-1}, \boldsymbol{x}_{t-1:t})$. A gated convolutional update then computes

$$\boldsymbol{g}_t = \sigma(f_g[Z_t, \bar{\boldsymbol{m}}_{t-1}, A_t, C_t])\,, \tag{3}$$

$$\boldsymbol{m}_t = \boldsymbol{g}_t \odot f_z(Z_t, C_t) + (1 - \boldsymbol{g}_t) \odot \bar{\boldsymbol{m}}_{t-1}, \tag{4}$$

where A_t is per-cell observation age and C_t is a sparse contact map. If $d_t = 0$, Z_t is masked and the update propagates memory while increasing age. Unlike a standard ConvGRU, the gate receives explicit evidence age, preventing unchanged hidden state from being interpreted as fresh evidence.

When foot j establishes a stable contact, forward kinematics yields $\tilde{h}_j$. We update nearby cells using a Gaussian kernel $k_j(i)$ and contact variance σ_c^2:

$$\hat{h}'_{t,i} = \frac{w_{t,i}\hat{h}_{t,i} + k_j(i)\sigma_c^{-2}\tilde{h}_j}{w_{t,i} + k_j(i)\sigma_c^{-2}}. \tag{5}$$

The same event reduces local risk only after slip and impact checks pass.

4.2 Calibrated Risk Head

Two decoder heads predict height $\hat{H}_t$ and risk logit L_t. Training sequences contain random rectangular occluders, dropped frames, motion blur, and depth-dependent noise. The risk label is one if the event in Eq. (1) occurs within the cell. The loss is

$$\mathcal{L} = \lambda_h \|M \odot (\hat{H} - H^\star)\|_1 + \lambda_n \mathcal{L}_{\text{normal}} \\ + \lambda_r\,\mathrm{BCE}(L, Y) + \lambda_t \|\hat{H}_t - W(\hat{H}_{t-1})\|_1, \tag{6}$$

where M marks valid ground truth. After training, a scalar temperature τ is fitted on a held-out validation set and $R = \sigma(L/\tau)$. We report expected calibration error (ECE) with 15 equal-mass bins and the Brier score; lower is better for both.

4.3 Risk-Constrained Foothold Control

The MPC stage extends a conventional centroidal objective with map terms:

$$J = J_{\text{track}} + \lambda_e J_{\text{effort}} + \sum_{k=1}^{K}\Big[\lambda_s \|\nabla \hat{H}(\boldsymbol{p}_k)\|_2^2 + \lambda_r\,\phi(r(\boldsymbol{p}_k))\Big]\,, \tag{7}$$

where $\phi(r) = -\log(1 - r + 10^{-4})$. Candidate contacts violating Eq. (2) are removed before continuous trajectory refinement. If no feasible candidate remains for three replans, the robot lowers its body, shortens step length, and turns toward the lowest-risk visible sector. This deterministic fallback avoids asking the learned map to hallucinate a route through wholly unobserved space.

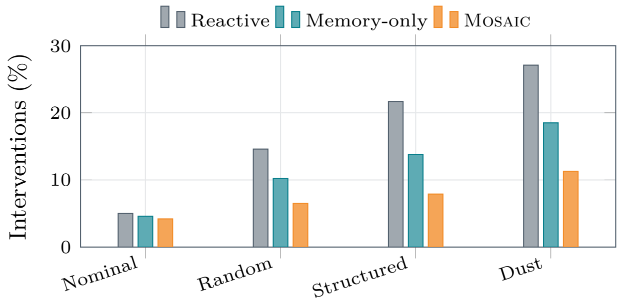

Figure 2: Illustrative hardware intervention rate across sensing conditions. Structured dropout removes every fourth 0.8-s camera interval; dust uses a scattering overlay plus real missing-depth masks.

5 Experimental Design

5.1 Platforms and Environments

Simulation uses a 12-DoF, 32-kg quadruped in 1,200 procedurally generated 12×8 m courses. Obstacles include steps (0.08–0.28 m), gaps (0.10–0.35 m), ramps (5–28°), and compliant patches. Dynamics randomize mass ($\pm 12\%$), friction (0.35–1.1), latency (0–35 ms), and camera extrinsics ($\pm 1.5°$).

Hardware evaluation uses the same controller on a 31-kg electric quadruped with a forward stereo camera at 30 Hz, state estimation at 400 Hz, and MPC at 50 Hz. Twelve physical courses combine pallets, foam blocks, gravel, and a three-step staircase. Each of four sensing conditions is repeated once per course, yielding 48 trials and 2.7 km of travel. An intervention is counted when a safety operator touches the stop control or supports the robot. Trials end after a fall, two interventions, or the 25-m goal.

5.2 Baselines and Metrics

We compare: (i) **Reactive**, which uses only the latest elevation frame; (ii) **DecayMap**, an exponential temporal filter whose confidence decays with age; (iii) **Memory-only**, the recurrent estimator without the risk head; and (iv) Mosaic. All methods share locomotion tracking, state estimation, candidate footholds, and compute budget. Primary metrics are course success and operator interventions. We also report velocity normalized by the no-dropout command, height mean absolute error (MAE), and risk calibration.

5.3 Implementation Details

The memory encoder has four residual blocks with channels $(24, 32, 48, 64)$; the recurrent state has 64 channels at quarter resolution and 1.8 million parameters. We train for 180k steps with AdamW, batch size 16, initial learning rate 3×10^{-4}, and cosine decay. Loss weights $(\lambda_h, \lambda_n, \lambda_r, \lambda_t) = (1.0, 0.25, 0.6, 0.15)$ are selected once on validation data. The map spans 6.4×4.8 m. Estimation takes 5.8 ms on an embedded GPU; candidate evaluation and MPC take 4.1 and 7.6 ms, respectively.

6 Results

6.1 Navigation Reliability

Table 1 shows that geometry memory alone improves simulated success by 13.4 points over Reactive, confirming the value of evidence across occlusions. Adding calibrated risk produces a further 7.3-point gain while reducing normalized speed by only 0.02 relative to Memory-only. On hardware, Mosaic completes 44 of 48 trials. Its four failures consist of two correlated map-warp errors, one unmodeled compliant obstacle, and one state-estimator reset. Figure 2 shows the largest benefit under structured dropout, where old terrain can appear plausible for several steps.

A paired bootstrap over courses gives a -6.1 percentage-point intervention difference between Mosaic and Memory-only (95% interval $[-10.4, -2.2]$). The result suggests that better reconstruction alone does not explain reliability: exposing the residual error to the controller changes which footholds are selected.

6.2 Calibration and Selective Planning

Before temperature scaling, the risk head has ECE 0.091 and is overconfident on cells older than 0.6 s. Scaling reduces ECE to 0.036 and Brier score from 0.118 to 0.082. When the planner rejects the riskiest 20% of candidate footholds, the remaining set contains 91% of valid contacts and only 24% of invalid contacts. This selective behavior matters more than raw height MAE: two methods can reconstruct similar surfaces while assigning very different costs to unobserved gap edges.

6.3 Ablations

Table 2 separates representation and control choices. Removing temporal memory causes the largest success loss because the planner repeatedly encounters empty foothold sets. Removing calibration causes the highest speed but nearly doubles interventions: unscaled scores are too confident under long occlusions. Contact correction primarily helps after a foot lands on gravel or foam, when visual height is biased. The fallback affects only 8.3% of replans yet prevents five falls, indicating that explicit behavior for “nothing is trustworthy” is valuable.

7 Discussion

The core finding is not that a recurrent map always predicts geometry more accurately; it is that map confidence must have semantics that the planner can use. Three design rules emerge. First, age is evidence, not merely a feature: it should evolve even when perception is absent. Second, contact supplies sparse but decision-relevant ground truth and should update both geometry and trust. Third, calibration must be evaluated under the exact corruptions that matter at deployment, including bursts rather than only independent frame loss.

The risk budget δ also offers a practical operating knob. Reducing it from 0.20 to 0.08 lowers interventions by 2.4 points but reduces speed from 0.92 to 0.84; increasing it to 0.35 recovers nominal speed but erases half of the reliability gain. A deployment can therefore select a doc-

Table 1: Illustrative aggregate results. Mean ± standard error over simulation seeds; hardware success is reported over 48 trials. The best value is bold and the second best is underlined. All planners use the same locomotion controller and sensing stream.

Method	Sim. success ↑	HW success ↑	Interventions ↓	Speed ratio ↑	Height MAE (cm) ↓	ECE ↓
Reactive	71.4±1.8	31/48	17.1%	**0.97**	7.8	—
DecayMap	79.6±1.5	35/48	13.5%	0.90	6.1	0.142
Memory-only	84.8±1.3	39/48	10.9%	<u>0.94</u>	<u>4.7</u>	—
MOSAIC	92.1±0.9	**44/48**	**7.5%**	0.92	**4.1**	**0.036**

Table 2: Ablation under structured dropout. "Contact" applies Eq. (5); "cal." applies temperature calibration.

Variant	Success ↑	Interv. ↓	Speed ↑
Full MOSAIC	**90.8**	**7.9**	0.92
− contact correction	86.1	10.7	0.91
− calibrated risk	82.9	13.2	**0.95**
− temporal memory	76.8	16.9	0.93
− fallback behavior	84.0	12.6	0.94

umented safety–throughput tradeoff without retraining locomotion.

7.1 Limitations and Responsible Use

The 2.5D map cannot represent overhangs, rails, or multiple surfaces in one cell. Equation (2) also ignores spatially correlated errors, so its probability is not a formal system-level safety guarantee. Fast deformable terrain can invalidate contact correction, and the fallback assumes enough room to lower and rotate the body. Finally, the reported study is an illustrative template rather than evidence for a deployable product. Real evaluation should publish seeds and code, include diverse weather and terrain, report all safety interventions, and retain an independent emergency stop. The method should not be used around people without a separate certified safety layer.

8 Conclusion

We presented MOSAIC, a terrain-memory interface that estimates both local geometry and the probability that geometry will mislead a foothold decision. By propagating evidence, repairing it with contact, calibrating risk, and constraining accumulated foothold risk, the system remains useful during perceptual dropout instead of silently treating stale terrain as truth. The illustrative experiments show complementary gains from temporal memory and risk-aware control. Future work should model correlated map errors, extend the representation to full 3D structure, and couple locomotion with active viewpoint selection.

Acknowledgments

This sample manuscript uses fictional authors, affiliations, experiments, and numerical results for demonstration. It was designed as a complete, self-contained LaTeX starting point; no claim of real-world performance is intended.

References

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

[2] 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.

[3] Ruben Grandia, Fabian Jenelten, Shaohui Yang, Farbod Farshidian, and Marco Hutter. Perceptive locomotion through nonlinear model-predictive control. *IEEE Transactions on Robotics*, 39(5):3402–3421, 2023.

[4] Armin Hornung, Kai M. Wurm, Maren Bennewitz, Cyrill Stachniss, and Wolfram Burgard. OctoMap: An efficient probabilistic 3d mapping framework based on octrees. *Autonomous Robots*, 34(3):189–206, 2013.

[5] Helen Oleynikova, Zachary Taylor, Marius Fehr, Roland Siegwart, and Juan Nieto. Voxblox: Incremental 3d euclidean signed distance fields for on-board MAV planning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1366–1373, 2017.

[6] Péter Fankhauser, Marko Bjelonic, C. Dario Bellicoso, Takahiro Miki, and Marco Hutter. Robust rough-terrain locomotion with a quadrupedal robot. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 5761–5768, 2018.

[7] Nikita Rudin, David Hoeller, Philipp Reist, and Marco Hutter. Learning to walk in minutes using massively parallel deep reinforcement learning. In *Proceedings of the Conference on Robot Learning*, volume 164, pages 91–100, 2022.

[8] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *Advances in Neural Information Processing Systems*, volume 30, 2017.

[9] Lukas Schmid, Michael Pantic, Raghav Khanna, Lionel Ott, Roland Siegwart, and Juan Nieto. An efficient sampling-based method for online informative path planning in unknown environments. *IEEE Robotics and Automation Letters*, 5(2):1500–1507, 2020.

[10] Lorenz Wellhausen, Alexey Dosovitskiy, René Ranftl, Krzysztof Walas, Cesar Cadena, and Marco Hutter. Where should i look? optimized gaze control for whole-body collision avoidance in dynamic environments. *IEEE Robotics and Automation Letters*, 5(2):3050–3057, 2020.