Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Rendered-Egocentric Vision for Short-Horizon Body-Risk Prediction in Quadruped Command Selection

Yuxiang Tian

yt3019@columbia.edu

Columbia University

Research Article

Keywords: Quadruped robots, egocentric vision, risk prediction, command selection, legged locomotion

Posted Date: September 8th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10337586/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Rendered-Egocentric Vision for Short-Horizon Body-Risk Prediction in Quadruped Command Selection

YuxiangTian[1]*

Affiliation

[1]*Department of Mechanical Engineering, Columbia University, New York, NY, 10027, USA*

Corresponding author: Yuxiang Tian, yt3019@columbia.edu

Abstract--Quadruped robots require local safety decisions before a velocity command is sent to a locomotion controller. This study investigates whether a rendered egocentric camera observation and a candidate command can predict short-horizon body-risk outcomes in simulation. We introduce a command-conditioned visual risk model above a fixed Unitree Go2 velocity-command controller. The deployed selector evaluates a small candidate set using only the current rendered view and each candidate command; it does not use terrain scans, map metadata, outcome labels, or rollout lookahead. Execution-time body feedback complements the visual model for contact- and progress-dependent failures. In MuJoCo experiments, the proposed model improves best-command agreement from 0.297 to 0.484 and safer-side agreement from 0.359 to 0.594 relative to an earlier visual-risk model. In matched short-horizon executions, the selector reduces the failure rate from 0.521 to 0.406 and body-ground contact from 0.427 to 0.240, while retaining 70% of nominal forward progress. Closed-loop trials identify the boundary of this result: local risk ranking does not by itself provide reliable long-horizon terrain traversal under a fixed gait controller. The study therefore presents a simulation-based local command-selection layer and a quantitative analysis of its safety-progress tradeoff.

Index Terms--Quadruped robots, egocentric vision, risk prediction, command selection, legged locomotion.

1 Introduction

Learning-based legged systems can couple perception, navigation, and locomotion to traverse difficult environments. A complementary local problem is less often isolated: before executing a candidate velocity command, can a robot estimate whether that command is likely to produce a short-horizon body-risk event?

End-to-end safe locomotion and navigation systems jointly optimize perception, command generation, and recovery. For example, Agile But Safe learns a reach-avoid value network and recovery policy for collision-free quadruped locomotion [29], while Lee et al. integrate mobility-aware navigation with locomotion for wheeled-legged robots [30]. Traversability, costmap, and safe-planning methods also provide essential tools for risk-aware navigation [5]-[8]. This paper does not attempt to match those integrated navigation systems. Instead, it studies a narrower question: whether rendered egocentric vision and a candidate velocity command can provide a useful pre-action estimate of body-risk above a fixed gait controller.

We propose a simulation-based decision module that estimates short-horizon body risk from rendered egocentric vision and candidate velocity commands. The module does not replace the gait controller, provide a safety guarantee, or claim to solve long-horizon navigation. It biases local command selection toward lower-risk visible motion, while execution-time body feedback handles contact-dependent failures.

This study makes three scoped contributions. First, it formulates rendered-vision and command-conditioned body-risk prediction above a fixed quadruped gait controller. Second, it evaluates the resulting local selector with matched command-ranking, route-cognition, and ablation protocols. Third, it quantifies the short-horizon safety-progress tradeoff and uses closed-loop trials only to diagnose the remaining route-to-gait execution gap.

2 Related Work

Learning-based quadruped locomotion has advanced through large-scale reinforcement learning, adaptation, imitation, and vision-conditioned control [11]-[17]. Recent systems also integrate safety mechanisms and navigation with locomotion [29], [30]. These works improve a policy or a complete navigation stack. Our work uses a different and more limited decomposition: the low-level gait controller is fixed, while a separate module predicts short-horizon body risk for candidate commands before execution.

Traversability, costmap, and terrain-affordance methods estimate where a robot can travel from geometry, semantics, vision, or proprioceptive interaction [18]-[23]. These methods are closely related, but their decision variable is usually a map cell, terrain patch, or path segment. In contrast, our decision variable is a candidate command, and the predicted quantity is short-horizon body failure risk under that command.

Risk-aware autonomy has also been studied through formal risk measures, constrained planning, safe learning, recovery policies, and failure discovery [24]-[28]. Our work follows the same motivation of reducing physical failure, but uses a different operating point. Rather than learning a complete risk-sensitive locomotion policy, we learn a pre-action visual body-risk prior that biases local command selection before the robot moves.

3 Method

A. Overview

We evaluate the decision layer in a MuJoCo simulation of a Unitree Go2 quadruped. As shown in Fig. 1, the risk-aware module operates above a fixed velocity-command locomotion controller. At each decision step, the robot receives a rendered head-camera grid and evaluates ten candidate velocity commands: center forward, slow forward, crawl, left/right forward, hard left/right turns, fast forward,

and left/right sidesteps. The command set includes forward speeds from 0.018 to 0.140 m/s, lateral speeds up to 0.040 m/s in magnitude, and yaw rates up to 0.200 rad/s in magnitude. This setup isolates local command ranking from low-level gait generation.

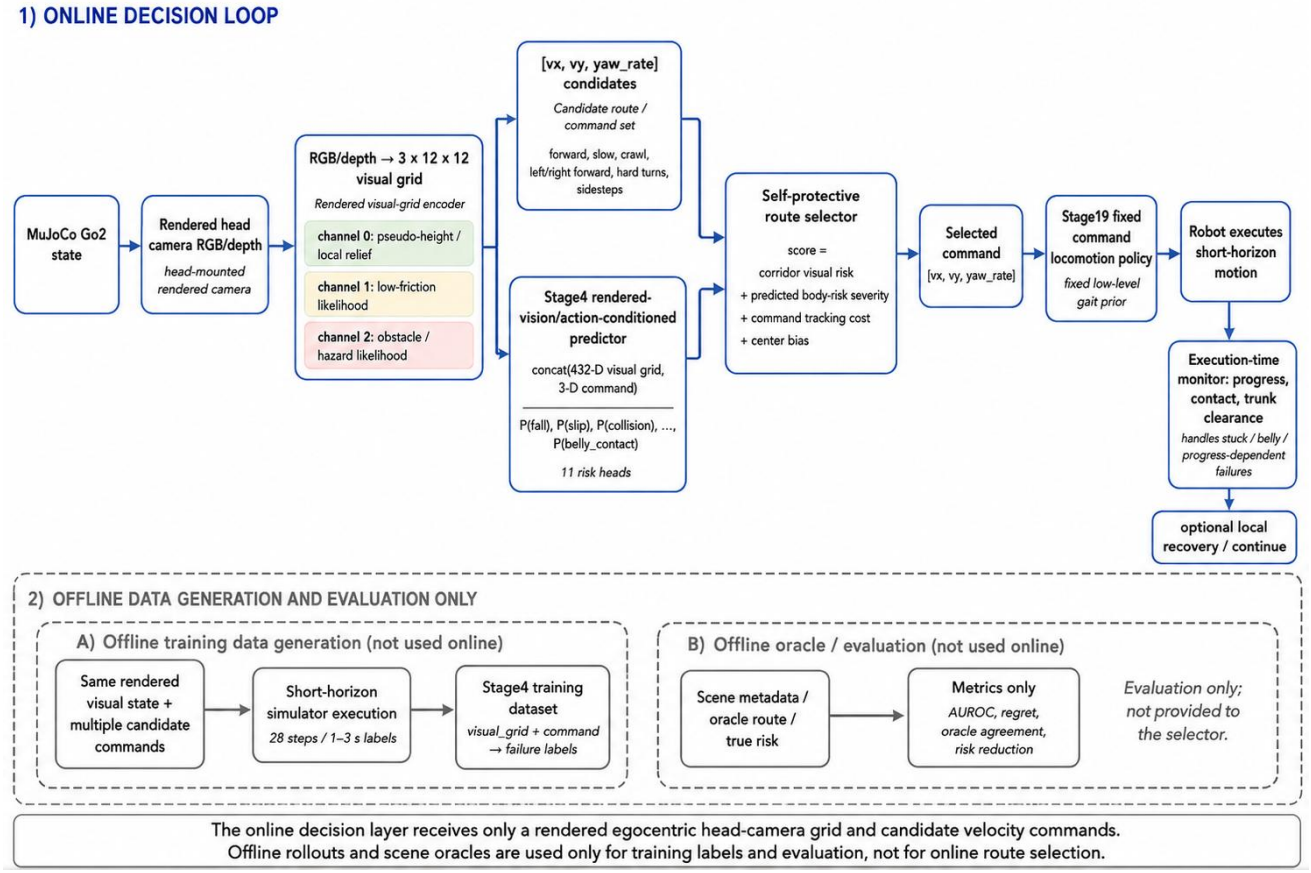


Fig. 1 System overview and local command-selection formulation

The evaluation uses two terrain families. The first consists of structured safe-corridor scenes, where a visibly lower-risk route exists within the camera view. These scenes test whether the decision layer can select a safer direction than nominal forward walking. The second consists of more cluttered relative-risk scenes, where no route is entirely safe, and the objective is to choose a less damaging direction. This setting is used to evaluate risk-aware decision-making under ambiguous terrain conditions, not to claim fully solved long-horizon navigation.

We compare the proposed action-conditioned visual risk model with four baselines: nominal forward walking, random visual-corridor selection, a handcrafted visual heuristic, and an earlier learned visual-risk model with weaker command conditioning. An oracle selector is used only as an offline upper bound for evaluation. We report results under three protocols: offline command-ranking evaluation, short-horizon embodied execution, and closed-loop diagnostic trials. The closed-loop trials are included to identify the remaining gap between high-level route intent and low-level gait execution.

B. Action-Conditioned Visual Body-Risk Prediction

The predictor uses a 432-dimensional rendered visual grid (three channels by 12 by 12 cells) concatenated with the candidate command c = [vx, vy, yaw_rate]. It is implemented as independent MLP risk heads with hidden layers of 128 and 64 units, standardization, early stopping, and one binary classifier per risk label. The training set contains 9,292 visual-command samples and 11 labels. Positive labels are obtained from 28-step short-horizon simulator execution after applying each candidate command. The command-sensitive subset deliberately pairs the same visual state with different commands. Terrain metadata, terrain scans, oracle outcomes, and online rollout lookahead are not predictor inputs.

The learned heads produce label-wise probabilities that are combined into a scalar body-risk score for command ranking. This score is a fixed severity-weighted index rather than an absolute probability of mission failure. The offline oracle is defined only after execution as the candidate with the lowest observed short-horizon total risk under the same visual state and candidate set. Thus, the deployed selector receives only the rendered visual grid and the candidate command, while evaluation labels and scene metadata remain post-hoc.

We compare the proposed model with an earlier visual-command risk model trained using weaker command-sensitive data. The comparison tests whether pairing the same visual scene with different candidate commands improves command-conditioned risk ranking.

C. Self-Protective Command and Route Selection

The proposed decision layer treats route selection as a pre-action risk assessment problem. At each decision step, the robot observes the terrain ahead from its egocentric camera and considers a small set of feasible motion hypotheses, including nominal forward motion and laterally biased alternatives. Each hypothesis corresponds to a short-horizon motion intention that can be executed by the underlying velocity-command controller.

For each motion hypothesis, the visual risk model estimates the likelihood that executing the corresponding command will lead to a body-risk event. The selector minimizes a fixed score

$$S(\mathbf{u}) = 2.5R_{\text{grid}}(\mathbf{u}) + R_{\text{pred}}(\mathbf{u}) + \| \mathbf{u} - \mathbf{u}_0 \|_2^2 + 0.12 \, | \, d(\mathbf{u}) \, | \, ,$$

where R_{grid} is the rendered-corridor risk, R_{pred} is the predicted body-risk severity, u_0 is the nominal forward command, and $d(\mathbf{u})$ is the lateral corridor offset.

This procedure does not require a global map or future simulation rollouts. It only uses the current egocentric observation and the command under consideration, making the decision layer compatible with a fixed locomotion controller.

The selected command is not intended to be an optimal long-horizon path plan. Instead, it represents a local self-protective choice: when the nominal forward direction appears risky, the robot can bias its motion toward a visually safer corridor. If all visible alternatives are risky or uncertain, the selected behavior may become more conservative, which can reduce physical failure but also reduce forward progress. This safety-progress tradeoff is evaluated explicitly in the embodied experiments.

Offline oracle labels are used only after execution or dataset evaluation to measure how close the selected command was to the lowest-risk available option. They are not provided to the selector during decision-making. This distinction is important because the goal is not to imitate an oracle planner, but to test whether a learned visual body-risk prior can support safer local command selection before the robot moves.

D. Execution Monitor and Recovery Interface

The visual risk model estimates danger before the robot moves, but several failure modes cannot be reliably resolved from vision alone. Loss of progress, body-ground contact, repeated foot slippage, and unstable posture often become clear only after the commanded motion is executed. We therefore use execution-time feedback as a complementary safeguard rather than treating the visual predictor as a complete failure detector.

During execution, the robot monitors whether realized motion remains consistent with the selected command. The short-horizon monitor flags stuck behavior when the command speed exceeds 0.045 and forward progress over a 20-step window is below 0.004 m. It flags belly-contact risk when trunk clearance is below 0.18 m for at least two consecutive steps. When such evidence appears, the system can interrupt command following and switch to a local recovery behavior.

This interface also defines the boundary of the present system. The recovery layer is local and reactive; it is not intended to solve global route planning or guarantee completion in highly contact-rich terrain. As a result, the closed-loop experiments are interpreted as a test of how well high-level risk-aware route intent can be realized by the underlying gait controller. When the robot identifies a safer direction but cannot physically track it with sufficient lateral or yaw authority, the failure is attributed to the route-to-gait execution interface rather than to the visual risk estimate alone.

4 EXPERIMENTS

A. Benchmarks and Baselines

Experiments were conducted using a simulated Unitree Go2 quadruped in MuJoCo. The evaluation includes four 10 m by 10 m structured safe-corridor scenes and four relative-risk Complex V2 scenes. The safe-corridor scenes test whether the decision layer moves away from risky nominal forward motion. The Complex V2 scenes test whether it selects a less damaging local direction when every visible route carries some risk.

All methods use the same fixed low-level command controller and the same candidate motion set. During decision-making, the selector only receives the egocentric visual observation and candidate command. Terrain labels, outcome labels, and offline risk references are used only for evaluation.

We compare the proposed visual-command risk selector with four baselines. The nominal baseline always commands forward motion and represents the fixed locomotion controller without visual risk reasoning. The random corridor baseline selects among visible alternatives without learned risk assessment. The geometric heuristic baseline uses handcrafted visual cues to avoid apparently unsafe regions. The earlier learned baseline uses the same visual-command interface but weaker command-sensitive training. An offline reference selector is additionally reported as an upper bound, but it is not available to any deployed method.

B. Metrics

We use three groups of metrics to evaluate prediction quality, decision quality, and embodied performance. Prediction quality is measured by held-out discrimination and calibration, including AUROC, false-negative rate, expected calibration error, and Brier-type probabilistic scores. These metrics evaluate the model as a pre-action risk prior rather than as a post-failure detector.

For decision quality, we compare the selected motion with an offline reference computed from observed outcomes. Top-1 agreement means selecting the same candidate as the oracle, whereas side agreement ignores speed and checks whether the selected lateral direction matches the oracle. True-risk regret is defined as $R(\mathbf{u}_{sel}) - R(\mathbf{u}^*)$, where $\mathbf{u}^*$ is the offline lowest-risk candidate. Risk reduction relative to nominal forward motion is defined as $R(\mathbf{u}_{center}) - R(\mathbf{u}_{sel})$.

These metrics are used only for evaluation; the offline reference is never provided to the robot during decision-making.

For embodied performance, we measure physical safety together with task progress. Safety is evaluated by falls, slips, collisions, stuck events, belly contact, and severe event density. Progress is evaluated by retained forward motion and risk-adjusted progress, where retained progress is normalized by center-forward progress on the same sampled states and risk-adjusted progress is defined as $\frac{\max(p, 0)}{1 + N_{\text{core}}}$, where p is forward progress and N_{core} is the number of core risk events.

We also report completion and event density in closed-loop trials to distinguish local safety improvement from full route completion. This combination is necessary because a self-protective policy is only useful if it reduces risk without simply stopping the robot.

5 RESULTS

A. Predictor Calibration and Failure-Modality Boundary

The first question is whether egocentric vision can provide a useful estimate of body risk before the robot moves. As shown in Fig. 2, the proposed predictor gives informative pre-action risk estimates, but reliability depends strongly on the failure type. The aggregate body-risk score provides a usable safety prior, with held-out discrimination around the moderate range for any-core failure. This supports its use for ranking candidate commands, rather than as a complete guarantee of safe execution.

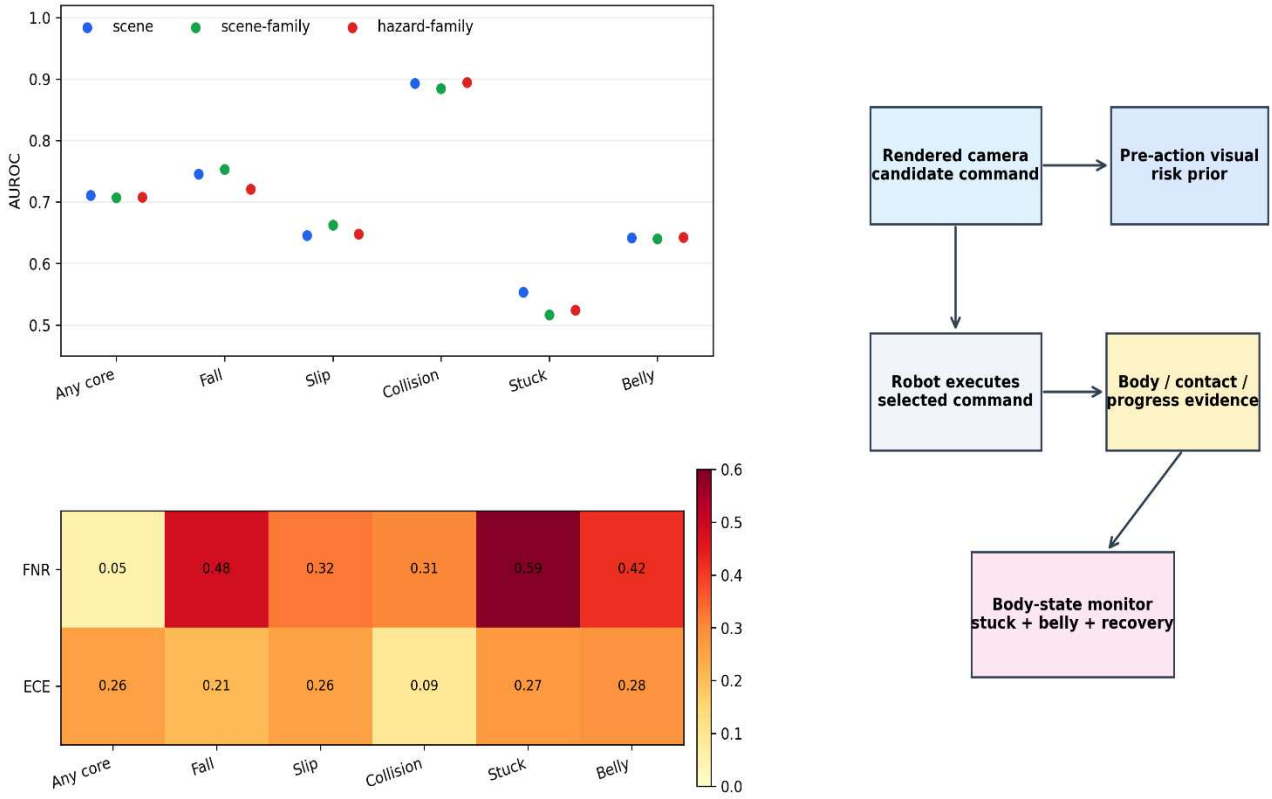


Fig. 2 Calibration results and the boundary between pre-action visual risk and execution-time body-state evidence

The strongest visual signal appears in collision-related failures. This is expected because obstacles and blocked corridors are often directly visible before motion. Falling and slipping also show partial visual predictability, since terrain roughness, slope, and surface appearance can indicate unstable execution conditions. However, these risks are less visually separable than collision because they also depend on contact dynamics and the realized gait response.

The weakest pre-action visual signal is associated with failures that emerge only during execution, such as loss of progress, local immobilization, and undesired body-ground contact. These cases define a modality boundary rather than a simple predictor failure. They are not purely perceptual events; they require feedback from body motion, contact state, and realized progress after the command has been executed. This result motivates the execution monitor introduced in Section III-D and explains why the proposed system combines pre-action visual risk prediction with reactive body-state supervision.

B. Command Ranking and Rendered-Vision Route Choice

The proposed model improves command ranking before execution. As summarized in Table I, it selects commands closer to the offline lowest-risk reference than the earlier model on the same visual states and candidate sets. Best-command agreement increases from 0.297 to 0.484, while safer-side agreement increases from 0.359 to 0.594. Average excess risk decreases from 0.069 to 0.025, and risk reduction relative to nominal forward walking increases from 0.040 to 0.084. These results show that the same visual scene can imply different short-horizon body risk when the intended command changes.

Table 1. Command-ranking and rendered-vision route-choice performance.

Metric	Earlier model	Proposed model	Delta	Direction
TOP-1 oracle agreement	0.297	0.484	0.187	higher
Side agreement	0.359	0.594	0.235	higher
Top-3 agreement	0.688	0.719	0.031	higher
Mean regret	0.069	0.025	-0.044	lower
Risk reduction vs center	0.040	0.084	0.044	higher
Risk reduction vs heuristic	0.034	0.078	0.044	higher
Unsafe selection rate	0.261	0.217	-0.044	lower

We next evaluate whether this command-level improvement produces meaningful route choices from egocentric vision. Fig. 3 shows representative examples in which the robot selects a visible corridor with a lower expected body risk than that of nominal forward walking. In structured corridor scenes, the selected direction often avoids visibly hazardous regions, resulting in a clear reduction in route risk. In cluttered relative-risk scenes, the task is harder because no visible direction is fully safe. In this setting, the selector is better interpreted as choosing a less damaging local direction rather than finding a globally safe path.

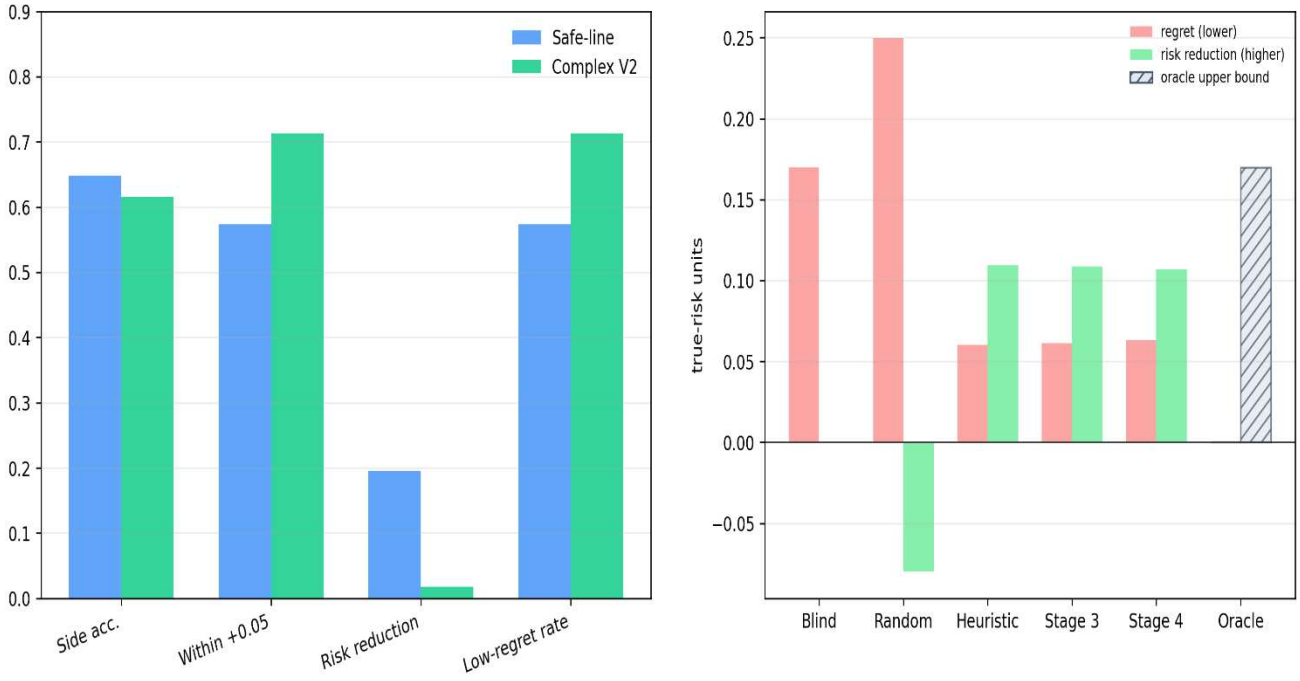


Fig. 3 Examples of command-conditioned route choice from rendered egocentric vision. Oracle routes are evaluation-only

The route-level results are therefore more nuanced than the command-ranking results. The learned selector improves over blind forward walking and arbitrary corridor selection, but it does not dominate a strong geometric heuristic in every scene. This indicates that part of the route-choice problem is already captured by visible terrain geometry. The learned predictor is most useful when different commands through the same visual scene lead to different body outcomes. Thus, the main benefit of the proposed layer is not oracle imitation, but risk-aware local command selection that biases the robot away from visually dangerous motion when safer alternatives are available.

C. Short-Horizon Execution and Closed-Loop Limitation

We next test whether safer command selection leads to safer embodied motion. As summarized in Fig. 4 and Table II, all methods are evaluated from the same 96 sampled short-horizon states using the same low-level controller, and the proposed selector reduces body-risk events compared with nominal forward walking. The overall failure rate decreases from 0.521 to 0.406, and undesired body-ground contact decreases from 0.427 to 0.240. The selected motions also achieve higher risk-adjusted progress, indicating that the improvement is not limited to offline ranking but transfers to physical execution through the quadruped controller.

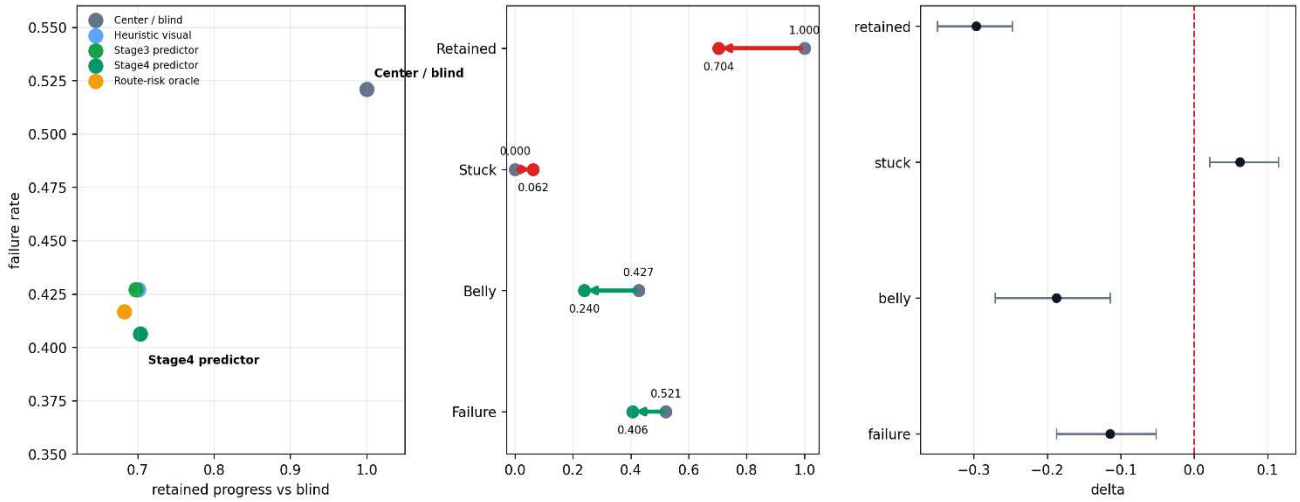


Fig. 4 Short-horizon execution safety-progress tradeoff under the fixed locomotion controller

Table 2. Short-horizon embodied execution safety-progress results.

Method	N	Failure	Progress m	Retained	Risk-adj.	Side acc.
Center/blind	96	0.521	0.146	1.000	0.028	0.365
Heuristic visual	96	0.427	0.102	0.701	0.037	0.604
Earlier model	96	0.427	0.102	0.698	0.036	0.594
Proposed predictor	96	0.406	0.103	0.704	0.039	0.583

Oracle	96	0.417	0.100	0.682	0.040	1.000

This safety gain comes with a clear progress tradeoff. The robot retains approximately 70% of the forward progress achieved by nominal walking, and some trials introduce local immobilization when the selected motion becomes overly conservative. This behavior is consistent with the objective of self-protective command selection: the robot reduces exposure to high-risk terrain, but it may sacrifice speed and directness when the visible alternatives are uncertain. The result should therefore be interpreted as safer short-horizon execution, not as uniformly better locomotion.

The closed-loop trials are small-N diagnostics rather than robust traversal validation, and they clarify the boundary of the current system. In structured corridor scenes, the full stack can be instantiated in a closed loop, showing that the visual decision layer, command interface, and execution monitor can operate together. In the more cluttered relative-risk scenes, however, full traversal remains unsolved. Adding local recovery reduces the density of severe body-risk events substantially, but it does not restore route completion. This indicates that the remaining bottleneck is the execution interface: the high-level selector may identify a safer local direction, but the low-level gait controller does not always provide enough lateral and yaw authority to realize that route intent over long horizons.

Overall, these results support a limited but useful claim. The proposed layer improves short-horizon physical safety and exposes where visual risk reasoning must be coupled with stronger route-following locomotion. The result is not a complete solution to complex-terrain navigation. It is a decision-level safety module whose benefit is strongest before and during local command execution.

6 Discussion and Conclusion

The results show that pre-action visual risk prediction can support self-protective decision making for quadruped locomotion, but only within a clearly defined scope. The proposed layer improves local command selection by estimating how different candidate motions may affect body safety before execution. This is most effective when risk is visually expressed, such as in blocked corridors, rough terrain regions, or visually unfavorable surface conditions. In these cases, the robot can bias its motion away from nominal forward walking and select a safer visible direction.

The results also show that visual risk prediction is not a complete substitute for execution-time feedback. Some failures become apparent only after contact is made, such as loss of progress, repeated slipping, undesired body-ground contact, or the inability to realize a selected route. These cases motivate a hybrid view of self-protection: vision can provide a useful prior before motion, while body-state feedback is needed to detect and respond to failures that emerge during execution.

A central observation is the difference between local decision quality and long-horizon route execution. The proposed model improves command ranking and short-horizon safety, but this does not automatically imply solved complex-terrain traversal. In more difficult scenes, the high-level selector may identify a safer local direction, yet the low-level gait controller may lack sufficient lateral or yaw authority to track that intent over time. This mismatch explains why local recovery can reduce severe body-risk events without guaranteeing route completion.

The present study has several limitations. All experiments are conducted in simulation, and the terrain set is controlled rather than exhaustive. The fixed locomotion controller isolates local command selection but limits lateral and turning authority. The recovery behaviors are local and reactive, not a substitute for integrated navigation. The reported oracles and risk labels are simulator-derived evaluation references, not inputs to the deployed selector. Consequently, the results should be interpreted as simulation evidence for local command ranking and short-horizon safety, not as validation of real-world or long-horizon autonomous navigation. Future work should evaluate the framework on physical quadrupeds, expand the terrain distribution, and jointly train risk-aware route selection with a route-following locomotion policy.

In conclusion, this paper studies a decision-level body-risk assessment layer for short-horizon quadruped command selection from rendered egocentric vision. The results show improved local command ranking and reduced short-horizon body-risk events under a fixed controller. They do not establish complete autonomous navigation in complex terrain. The remaining failures identify a clear direction for future work: predictive visual command ranking must be coupled with stronger route-following locomotion and physical-robot validation.

Funding Declaration: This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Competing interests: I have nothing to declare.

REFERENCES

[1] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter, "Learning agile and dynamic motor skills for legged robots," Science Robotics, vol. 4, no. 26, 2019.
[2] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning quadrupedal locomotion over challenging terrain," Science Robotics, vol. 5, no. 47, 2020.
[3] T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning robust perceptive locomotion for quadrupedal robots in the wild," Science Robotics, vol. 7, no. 62, 2022.
[4] D. Hoeller, N. Rudin, D. Sako, and M. Hutter, "ANYmal Parkour: Learning agile navigation for quadrupedal robots," Science Robotics, vol. 9, no. 88, 2024.
[5] S. Fahmi, V. Barasuol, D. Esteban, O. Villarreal, and C. Semini, "ViTAL: Vision-based terrain-aware locomotion for legged robots," IEEE Transactions on Robotics, vol. 39, no. 2, pp. 885-904, 2023.
[6] J. Frey, M. Mattamala, N. Chebrolu, C. Cadena, M. Fallon, and M. Hutter, "Fast traversability estimation for wild visual navigation," in Robotics: Science and Systems, 2023.
[7] R. Yang, M. Zhang, N. Hansen, H. Xu, and X. Wang, "Learning vision-guided quadrupedal locomotion end-to-end with cross-modal transformers," in International Conference on Learning Representations, 2022.
[8] A. Dixit, D. D. Fan, K. Otsu, S. Dey, A.-A. Agha-Mohammadi, and J. W. Burdick, "SPOMP: Stochastic traversability evaluation and planning for risk-aware navigation; results from the DARPA Subterranean Challenge," IEEE Transactions on Field Robotics, 2024.

[9] G. B. Margolis, X. Fu, Y. Ji, and P. Agrawal, "Learning to see physical properties with active sensing motor policies," in Conference on Robot Learning, 2023.
[10] A. Marco, D. Baumann, M. Khadiv, P. Hennig, L. Righetti, and S. Trimpe, "Robot learning with crash constraints," IEEE Robotics and Automation Letters, vol. 6, no. 2, pp. 1439-1446, 2021.
[11] N. Rudin, D. Hoeller, P. Reist, and M. Hutter, "Learning to walk in minutes using massively parallel deep reinforcement learning," in Proceedings of the 5th Conference on Robot Learning, PMLR, vol. 164, pp. 91-100, 2022.
[12] A. Kumar, Z. Fu, D. Pathak, and J. Malik, "RMA: Rapid motor adaptation for legged robots," in Robotics: Science and Systems, 2021.
[13] X. B. Peng, E. Coumans, T. Zhang, T.-W. E. Lee, J. Tan, and S. Levine, "Learning agile robotic locomotion skills by imitating animals," in Robotics: Science and Systems, 2020.
[14] L. M. Smith, J. C. Kew, X. B. Peng, S. Ha, J. Tan, and S. Levine, "Legged robots that keep on learning: Fine-tuning locomotion policies in the real world," in IEEE International Conference on Robotics and Automation, 2022.
[15] A. Agarwal, A. Kumar, J. Malik, and D. Pathak, "Legged locomotion in challenging terrains using egocentric vision," in Conference on Robot Learning, 2023.
[16] J. Jin, C. Zhang, J. Frey, N. Rudin, M. Mattamala, C. Cadena, and M. Hutter, "Resilient legged local navigation: Learning to traverse with compromised perception end-to-end," in IEEE International Conference on Robotics and Automation, 2024.
[17] C. Li, M. Vlastelica, S. Blaes, J. Frey, F. Grimminger, and G. Martius, "Learning agile skills via adversarial imitation of rough partial demonstrations," in Conference on Robot Learning, 2023.
[18] G. Kahn, P. Abbeel, and S. Levine, "BADGR: An autonomous self-supervised learning-based navigation system," IEEE Robotics and Automation Letters, vol. 6, no. 2, pp. 1312-1319, 2021.
[19] M. Mattamala, J. Frey, P. Libera, N. Chebrolu, G. Martius, C. Cadena, M. Hutter, and M. Fallon, "Learning-based visual navigation: Fast traversability learning via pre-trained models and online self-supervision," Autonomous Robots, vol. 49, art. no. 19, 2025.
[20] D. D. Fan, S. Dey, A.-A. Agha-Mohammadi, and E. A. Theodorou, "Learning risk-aware costmaps for traversability in challenging environments," IEEE Robotics and Automation Letters, vol. 7, no. 1, pp. 279-286, 2022.
[21] M. G. Castro, S. Triest, W. Wang, J. M. Gregory, F. Sanchez, J. G. Rogers III, and S. Scherer, "How does it feel? Self-supervised costmap learning for off-road vehicle traversability," in IEEE International Conference on Robotics and Automation, 2023.
[22] M. V. Gasparino, A. N. Sivakumar, Y. Liu, A. E. B. Velasquez, V. A. H. Higuti, J. Rogers, H. Tran, and G. Chowdhary, "WayFAST: Navigation with predictive traversability in the field," IEEE Robotics and Automation Letters, vol. 7, no. 4, pp. 10651-10658, 2022.
[23] M. Elnoor, A. J. Sathyamoorthy, K. Weerakoon, and D. Manocha, "ProNav: Proprioceptive traversability estimation for legged robot navigation in outdoor environments," IEEE Robotics and Automation Letters, 2024.
[24] A. Majumdar and M. Pavone, "How should a robot assess risk? Towards an axiomatic theory of risk in robotics," in Robotics Research, pp. 75-84, 2020.
[25] A. Hakobyan, G. C. Kim, and I. Yang, "Risk-aware motion planning and control using CVaR-constrained optimization," IEEE Robotics and Automation Letters, vol. 4, no. 4, pp. 3924-3931, 2019.
[26] B. Thananjeyan, A. Balakrishna, U. Rosolia, F. Li, R. McAllister, J. E. Gonzalez, S. Levine, F. Borrelli, and K. Goldberg, "Recovery RL: Safe reinforcement learning with learned recovery zones," in Conference on Robot Learning, 2020.
[27] C. Dawson and C. Fan, "A Bayesian approach to breaking things: Efficiently predicting and repairing failure modes via sampling," in Conference on Robot Learning, 2023.
[28] L. Schneider, J. Frey, T. Miki, and M. Hutter, "Learning risk-aware quadrupedal locomotion using distributional reinforcement learning," in IEEE International Conference on Robotics and Automation, 2024.
[29] T. He, C. Zhang, W. Xiao, G. He, C. Liu, and G. Shi, "Agile But Safe: Learning Collision-Free High-Speed Legged Locomotion," in Robotics: Science and Systems, 2024.
[30] J. Lee, M. Bjelonic, A. Reske, L. Wellhausen, T. Miki, and M. Hutter, "Learning Robust Autonomous Navigation and Locomotion for Wheeled-Legged Robots," Science Robotics, vol. 9, no. 89, 2024, doi: 10.1126/scirobotics.adi9641.