

APREBot: Active Perception System for Reflexive Evasion Robot

Zihao Xu^{1*}, Kuankuan Sima^{2*}, Junhao Deng³, Zixuan Zhuang⁴, Chunzheng Wang², Ce Hao^{1,5†},
and Jin Song Dong¹

Abstract—Reliable onboard perception is critical for quadruped robots navigating dynamic environments, where obstacles can emerge from any direction under strict reaction time constraints. Single-sensor systems face inherent limitations: LiDAR provides omnidirectional coverage but lacks rich texture information, while cameras capture high-resolution detail but suffer from restricted field of view. We introduce APREBot (Active Perception System for Reflexive Evasion Robot), a novel framework that integrates reflexive evasion with active hierarchical perception. APREBot strategically combines LiDAR-based omnidirectional scanning with camera-based active focusing, achieving comprehensive environmental awareness essential for agile obstacle avoidance in quadruped robots. We validate APREBot through extensive Sim2Real experiments on a quadruped platform, evaluating diverse obstacle types, trajectories, and approach directions. Our results demonstrate substantial improvements over strong baselines in both safety metrics and operational efficiency, highlighting APREBot’s potential for dependable autonomy in safety-critical scenarios. Paper homepage: <https://aprebot-2026.github.io/>.

I. INTRODUCTION

Legged robots are increasingly expected to operate in unstructured and dynamic environments, where reliable independent onboard perception is essential for autonomous navigation and interaction. Typical sensing modalities include LiDAR [1]–[4], which provides accurate range information over a wide field of view; RGB-D cameras [5], [6], which capture both appearance and depth within a restricted perspective; and event cameras [7], which deliver high temporal resolution suited to rapid changes in the scene. These onboard sensors support core perception tasks such as environment mapping, obstacle detection, and navigation.

However, perceiving dynamic obstacles remains particularly challenging for legged robots. Such obstacles can approach from arbitrary directions, including lateral or rear blind spots, which require continuous omnidirectional awareness. At the same time, high-speed motion imposes strict reaction time constraints, leaving only a very limited time window to detect and respond before a potential collision. Given these constraints, LiDAR provides global coverage but it cannot reliably segment objects and lacks semantic perception [8], cameras offer high-resolution details but their limited field of view prevents global perception [9], and event cameras

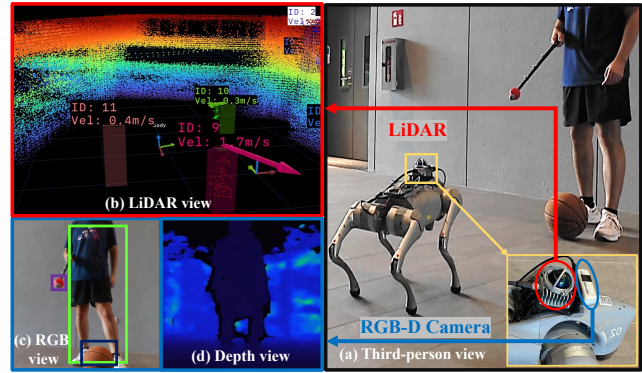


Fig. 1: Overview of APREBot’s active perception system. As illustrated in (a), the quadruped robot integrates LiDAR and RGB-D camera to realize active perception: as obstacles approach, the robot combines omnidirectional scanning with active focusing to enable threat-aware avoidance. The LiDAR-scanned point cloud in (b) provides omnidirectional perception, the RGB camera view in (c) enables high-resolution segmentation, and the depth camera view in (d) supports accurate 3D localization of the obstacle.

achieve fine temporal resolution but are sensitive to noise and lack robustness [10]. Therefore, no single sensor modality is sufficient, and combining them in a complementary manner is essential to achieve robust performance [11].

We propose **APREBot** (Active Perception System for Reflexive Evasion Robot), a novel framework that integrates reflexive evasion with active hierarchical perception to address dynamic obstacle avoidance under constrained reaction time. APREBot employs LiDAR-based omnidirectional scanning for global monitoring to ensure that obstacles approaching from arbitrary directions, including blind spots, can be reliably detected and tracked in real time. Once a candidate obstacle is identified, camera-driven active focusing is engaged to acquire high-resolution appearance details to distinguish high-risk obstacles (Fig. 1). A principled threat level formulation fuses these perception streams to assess obstacle severity and drive adaptive evasion. This perception-driven design builds on REBot [12], which demonstrated low-latency reflexive avoidance with privileged obstacle states, and extends it to real-world deployment without reliance on external devices.

We validate APREBot through extensive Sim2Real experiments on a quadruped platform, covering diverse dynamic obstacle conditions, including obstacles approaching from different directions, with varying speeds, and across multiple object types. The evaluation compares APREBot against strong perception baselines, highlighting the advantages of

* Equally contributed. † corresponding to cehao@u.nus.edu

¹ School of Computing, National University of Singapore, Singapore.

² Department of Electrical and Computer Engineering, National University of Singapore, Singapore.

³ School of Automation, Beijing Institute of Technology, China.

⁴ School of Computer Science and Engineering, Sun Yat-Sen University, China.

⁵ Beijing Zhongguancun Academy, China.

active hierarchical perception and threat-aware avoidance. Results demonstrate that APREBot consistently improves safety margins and reaction efficiency, while preserving the agility required for instantaneous evasive maneuvers.

In summary, the contributions of this paper are as follows:

- We propose APREBot, an active hierarchical perception framework that enables quadrupeds to perceive dynamic obstacles, evaluate their threat levels, and perform reflexive evasion without privileged information.
- We develop threat-aware perception mechanisms that combine LiDAR omnidirectional scanning, camera-based active tracking, short-horizon prediction, and a principled threat level for adaptive avoidance.
- We validate APREBot through extensive Sim2Real experiments, showing consistent improvements over representative perception baselines in safety, efficiency, and robustness.

II. RELATED WORKS

Single-sensor systems are widely explored for perception. Raw LiDAR methods, such as Omni-Perception [1], provide 360° geometric coverage but point clouds are inherently sparse, noisy, and distorted by motion, which hinders robust detection and tracking. Learning-based LiDAR detectors have advanced in autonomous driving [13], [14], yet suffer from domain gaps and poor cross-sensor generalization across platforms such as Avia, Velodyne, and MID-360. Vision-based approaches, including ABS [5] and dynamic interaction tasks such as badminton robots [15], provide semantically rich and smooth signals that support precise tracking. However, when applied to dynamic obstacle avoidance (DOA), LiDAR-only methods struggle with robust real-time detection and tracking [16], while vision-only methods remain constrained by their limited field of view and vulnerability to motion blur and occlusion.

Multi-sensor fusion has been investigated to overcome single-sensor limitations. Event-camera and frame-camera pipelines combine microsecond-level temporal resolution with semantic cues, improving robustness under rapid motion [7], but remain sensitive to illumination with limited semantic understanding. LiDAR-camera fusion approaches [17], [18], including IRS-based alignment [19], [20] and tightly coupled systems such as R2LIVE [21] and R3LIVE [22], enhance accuracy and robustness through joint optimization at the expense of computational efficiency. While recent works like LV-DOT [23] achieved real-time LiDAR-RGB fusion, existing methods still exhibit limitations in matching strategies, motion pattern filtering, and underutilize RGB-based 2D tracking robustness.

III. PRELIMINARY

Problem formulation. We study dynamic obstacle avoidance (DOA) for quadruped robots under a limited reaction time budget Δt_{react} . Each obstacle is modeled as a rigid body with radius r^O and states $\{p_t^O, v_t^O\}$ in the world frame \mathcal{W} . The robot (Unitree Go2) is a high-dimensional articulated system with base states $\{p_t^R, v_t^R, \omega_t^R, \theta_t^R\}$ and leg states $\{q_t^R, \dot{q}_t^R\}$ (12-DoF). Unless otherwise specified, all positions and velocities

TABLE I: Notation Summary

Symbol	Description
$p^R \in \mathbb{R}^3$	Robot base position (in \mathcal{W})
$v^R \in \mathbb{R}^3$	Robot base linear velocity (in \mathcal{W})
$\omega^R \in \mathbb{R}^3$	Robot base angular velocity (in \mathcal{W})
$\theta^R \in \mathbb{R}^3$	Robot base orientation (Euler angles, in \mathcal{B})
$q^R, \dot{q}^R \in \mathbb{R}^{12}$	Joint positions and velocities
$p^O, v^O \in \mathbb{R}^3$	Obstacle position and velocity (in \mathcal{W})
$r^O \in \mathbb{R}$	Obstacle radius
$\hat{p}_s^O, \hat{v}_s^O \in \mathbb{R}^3$	Estimated obstacle states from $s \in \{L, C\}$
$\hat{r}_s^O \in \mathbb{R}$	Estimated obstacle radius from $s \in \{L, C\}$
$id \in \mathbb{N}$	Obstacle track identity ($L \& C$)
$\Delta p^s \in \mathbb{R}^3$	Relative position, $\hat{p}_s^O - p^R$
$v_{\text{rel}}^s \in \mathbb{R}^3$	Relative velocity, $\hat{v}_s^O - v^R$
$\hat{d}^s \in \mathbb{R}^3$	Unit approach vector, $\Delta p^s / \ \Delta p^s\ $
$\mathcal{T}^L, \mathcal{T}^C \in \mathbb{R}_{>0}$	LiDAR / Camera threat levels
$r_{\text{safe}}, \alpha, \gamma, \epsilon \in \mathbb{R}_{>0}$	Proximity radius, weight, exponent, constant

World/body frames \mathcal{W}/\mathcal{B} . Time index t omitted when clear.

are expressed in the world frame \mathcal{W} ; orientation and joint states are expressed in the body-fixed frame \mathcal{B} . The complete summary of the notation is provided in Tab. I.

Active and hierarchical perception. APREBot relies solely on onboard sensing for obstacle perception. The LiDAR provides omnidirectional point clouds, which are segmented and tracked to obtain coarse obstacle estimates $\{\hat{p}_{t,L}^O, \hat{v}_{t,L}^O, \hat{r}_L^O, id\}$. Threat levels are computed for all detected obstacles, and the one with the largest LiDAR threat is used to adjust the robot's direction so that the corresponding obstacle enters the camera field of view. Once visible, the RGB-D camera provides refined estimates $\{\hat{p}_{t,C}^O, \hat{v}_{t,C}^O, \hat{r}_C^O, id\}$. This hierarchical perception pipeline exploits complementary sensing roles: LiDAR ensures 360° coverage with coarse estimates, while the camera contributes high-precision states for the most threatening obstacle.

Threat level. For a source $s \in \{L, C\}$ (LiDAR or Camera), we define a unified threat metric as in Eq. 1:

$$\mathcal{T}_t^s = \alpha \frac{\max(0, v_{\text{rel},t}^s \cdot \hat{d}_t^s)}{\|\Delta p_t^s\| + \epsilon} + (1 - \alpha) \left(\frac{r_{\text{safe}}}{\|\Delta p_t^s\| + \epsilon} \right)^\gamma, \quad (1)$$

where $\Delta p_t^s = \hat{p}_{t,s}^O - p_t^R$ is the relative position, $v_{\text{rel},t}^s = \hat{v}_{t,s}^O - v_t^R$ is the relative velocity, and $\hat{d}_t^s = \Delta p_t^s / \|\Delta p_t^s\|$ is the unit approach direction. The first term captures approach speed normalized by distance, while the second provides a proximity prior that yields nonzero threat for slow but nearby obstacles. Parameters $\alpha \in [0, 1]$, $\gamma \geq 1$, and r_{safe} control the balance between these two effects, and $\epsilon > 0$ avoids division by zero. LiDAR-based \mathcal{T}_t^L is always available; when the camera is engaged, \mathcal{T}_t^C is used for the refined target. In multi-obstacle scenarios, the maximum threat across obstacles is selected to indicate the most urgent direction. Compared with the classical time-to-collision (TTC), Eq. (1) offers greater robustness by remaining well-defined for slow or static yet proximate obstacles, while enabling a principled trade-off between distance- and velocity-based risk components.

IV. METHOD

In this section, we present APREBot, a hierarchical perception and threat-aware avoidance framework for dynamic

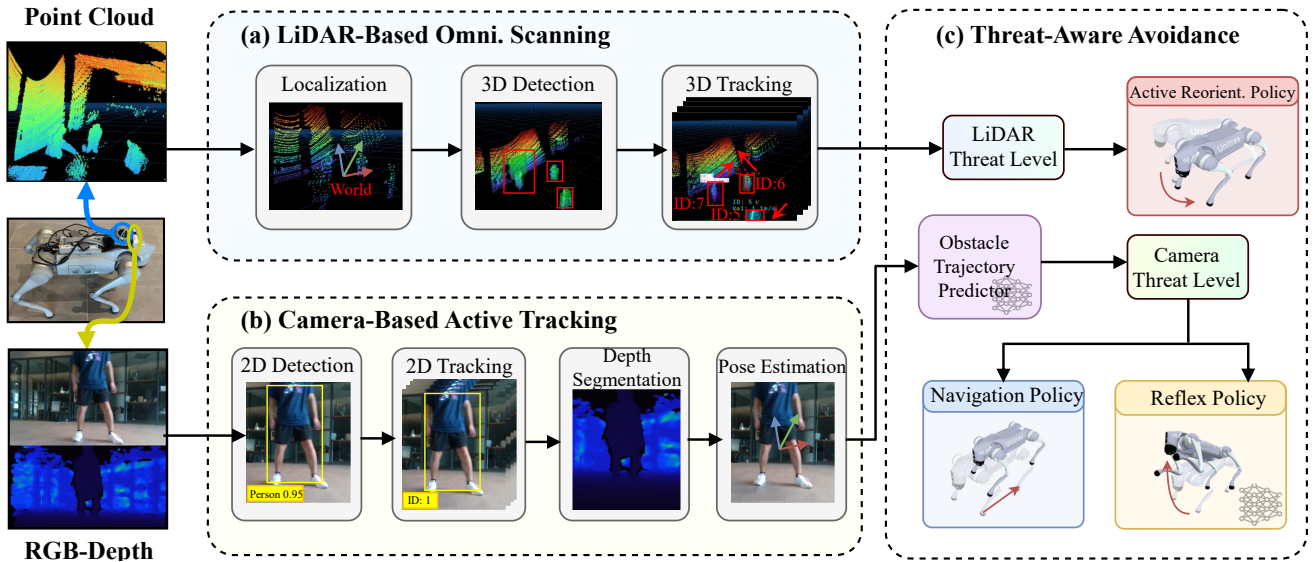


Fig. 2: Overview of the APREBot framework consisting of three stages. **(a) Omnidirectional Scanning Stage:** LiDAR-based perception operates on point cloud to perform localization and provides surrounding obstacle states through 3D detection and tracking. **(b) Active Tracking Stage:** The robot focuses on the most threatening obstacle, where the onboard RGB-D camera performs high-resolution 2D detection and tracking for segmentation, and uses depth-based refinement for accurate 3D state estimation. **(c) Threat-Aware Avoidance Stage:** Based on the evaluated threat level, the robot executes appropriate strategies, including active reorientation, navigation-based avoidance, and reflexive evasion.

obstacle avoidance. We outline the system structure (Sec. IV-A) with three stages—omnidirectional scanning, active tracking, and threat-aware avoidance—coordinated by a unified threat metric; describe hierarchical perception (Sec. IV-B); and the threat-aware avoidance policy (Sec. IV-C).

A. System structure

APREBot executes a three-stage perception–control cycle under the reaction time budget: Omnidirectional Scanning Stage, Active Tracking Stage, and Threat-Aware Avoidance Stage (Fig. 2).

(a) Omnidirectional scanning stage. The onboard LiDAR continuously segments and tracks obstacles, yielding coarse estimates $\{\hat{p}_{t,L}^O, \hat{v}_{t,L}^O, \hat{r}_L^O, id\}$ together with obstacle threats \mathcal{T}_t^L . The target is identified as the obstacle with the largest \mathcal{T}_t^L , which establishes a global priority while preserving omnidirectional situational awareness. Target identity is switched only if another obstacle keeps a higher threat for several consecutive steps, avoiding frequent changes in cluttered scenes; the approach direction $\hat{d}_t^L = \Delta p_t^L / \|\Delta p_t^L\|$ provides an orientation signal for subsequent acquisition.

(b) Active tracking stage. The robot adjusts its base orientation θ_t^R (and angular velocity ω_t^R) so that the active target enters the RGB-D camera field of view. Once visible, the camera provides refined segmentation of different objects (e.g., humans and balls) and estimates $\{\hat{p}_{t,C}^O, \hat{v}_{t,C}^O, \hat{r}_C^O, id\}$, with the corresponding threat \mathcal{T}_t^C replacing the coarse LiDAR quantities as primary inputs to decision-making. LiDAR continues to furnish panoramic coverage; if a newly detected obstacle exhibits a consistently larger \mathcal{T}_t^L than the current target, the perception focus is reassigned.

(c) Threat-aware avoidance stage. The entire avoidance process is driven by threat level. LiDAR first identifies

the obstacle with the highest threat and triggers active reorientation to bring it into the camera field of view. The camera then refines the threat-level assessment: lower levels lead to navigation-based retreat, while higher levels trigger reflexive evasions such as jumping or crouching.

This three-stage design preserves omnidirectional awareness, allocates selective refinement to the most critical target, and yields avoidance behaviors that adapt to Eq. (1) under the timing constraint Δt_{react} .

B. Hierarchical perception

APREBot employs a hierarchical perception system that combines the complementary strengths of LiDAR and cameras. LiDAR provides omnidirectional coverage, enabling reliable detection and tracking of obstacles approaching from any direction. The camera is then actively oriented toward selected targets to perform recognition and depth-based localization, allowing the system to distinguish obstacle types and refine their positions.

LiDAR-based omnidirectional scanning. Our LiDAR module follows a geometry-driven pipeline, avoiding reliance on domain-specific training data and thus remaining robust across heterogeneous LiDARs and unstructured environments. As shown in Alg. 1, the raw point cloud \mathcal{P}_t is first registered to the world frame using FAST-LIO, then downsampled and cropped by region-of-interest filters to reduce noise and computational load. Then, DBSCAN is applied to extract obstacle clusters, each represented by a 3D bounding box.

To maintain temporal consistency, bounding boxes across frames are associated through a Hungarian matcher with increased weight on horizontal displacement, reflecting the dominant motion pattern of approaching obstacles. Track identities are updated through a birth–death mechanism, while

Algorithm 1 LiDAR Perception

Input: LiDAR point cloud \mathcal{P}_t **Output:** Estimated obstacle states $\{\hat{p}_L^O, \hat{v}_L^O, \hat{r}_L^O, id\}$

```
1:  $\mathcal{P}_t^{world} \leftarrow \text{FastLIO}(\mathcal{P}_t)$ 
2:  $\mathcal{P}_t^{roi} \leftarrow \text{Filtering}(\mathcal{P}_t^{world})$ 
3:  $\mathcal{C}_t \leftarrow \text{DBSCAN}(\mathcal{P}_t^{roi})$ 
4:  $\mathcal{B}_t^{3D} \leftarrow \text{ComputeBBox}(\mathcal{C}_t)$ 
5: /* Hungarian Matching */
6: for  $i, j$  in  $\mathcal{B}_{t-1}^{3D} \times \mathcal{B}_t^{3D}$  do
7:    $C_{ij} \leftarrow 2\|\tilde{p}_{i,xy}^{t-1} - \tilde{p}_{j,xy}^t\|_2 + \|\tilde{p}_{i,z}^{t-1} - \tilde{p}_{j,z}^t\| - \text{IoU}_{ij}$ 
8:  $\mathcal{M}_t \leftarrow \text{Hungarian}(C)$ 
9:  $\mathcal{B}_t^{tracked} \leftarrow \text{AssignIDs}(\mathcal{B}_t^{3D}, \mathcal{M}_t)$ 
10: /* Motion Pattern Filtering */
11: for track  $b_i \in \mathcal{B}_t^{tracked}$  do
12:    $\Sigma_p \leftarrow \text{Cov}(\{\tilde{p}_i^{t-k:t}\})$ 
13:    $\sigma_v \leftarrow \text{Std}(\{\|v_i^{t-k:t}\|\})$ 
14:   if  $\text{tr}(\Sigma_p) > \epsilon_p$  OR  $\sigma_v > \epsilon_v$  then
15:     Keep track  $b_i$ 
16: /* Kalman Smoothing */
17:  $\{\hat{p}_L^O, \hat{v}_L^O, \hat{r}_L^O, id\} \leftarrow \text{KalmanFilter3D}(\mathcal{B}_t^{filtered})$ 
18: return  $\{\hat{p}_L^O, \hat{v}_L^O, \hat{r}_L^O, id\}$ 
```

Algorithm 2 RGB-Depth Perception

Input: RGB image \mathcal{I}_t^{rgb} , Depth image \mathcal{I}_t^{depth} , Camera params $\mathbf{K}, [\mathbf{R}|\mathbf{t}]$ **Output:** Estimated obstacle states $\{\hat{p}_C^O, \hat{v}_C^O, \hat{r}_C^O, id\}$

```
1: /* 2D Detection and Tracking */
2:  $\mathcal{D}_t^{2D} \leftarrow \text{YOLO}(\mathcal{I}_t^{rgb})$ 
3:  $\mathcal{D}_t^{tracked} \leftarrow \text{ByteTrack}(\mathcal{D}_t^{2D})$ 
4: /* BFS Instance Segmentation */
5: for detection  $d_i \in \mathcal{D}_t^{tracked}$  do
6:    $c_i \leftarrow \text{BBoxCenter}(d_i)$ 
7:    $\mathcal{M}_i \leftarrow \text{BFS}(c_i, \mathcal{I}_t^{depth}, \tau_{depth})$ 
8: /* 3D Pose Estimation */
9:  $\mathcal{P}_i^{cam} \leftarrow \{\mathbf{K}^{-1}[u \cdot z, v \cdot z, z]^T : (u, v) \in \mathcal{M}_i, z = \mathcal{I}_t^{depth}(u, v)\}$ 
10:  $\tilde{p}_i^{cam} \leftarrow \text{Centroid}(\mathcal{P}_i^{cam})$ 
11:  $\hat{p}_{C,i}^O \leftarrow \mathbf{R}\tilde{p}_i^{cam} + \mathbf{t}$ 
12:  $\hat{r}_{C,i}^O \leftarrow \text{EstimateRadius}(\mathcal{P}_i^{cam})$ 
13: /* Velocity Estimation */
14:  $\hat{v}_{C,i}^O \leftarrow \text{Diff}(\hat{p}_{C,i}^O)$ 
15: return  $\{\hat{p}_C^O, \hat{v}_C^O, \hat{r}_C^O, id\}$ 
```

a motion-consistency filter suppresses spurious tracks that fail position or velocity stability tests. Finally, each valid track is smoothed with a 3D Kalman filter, yielding obstacle states $(\hat{p}_L^O, \hat{v}_L^O, \hat{r}_L^O, id)$ that are maintained online for subsequent threat evaluation and avoidance control.

RGB-D-based active focusing. When an active target enters the camera’s field of view, an RGB-D pipeline is activated to provide precise and metrically consistent obstacle states. As summarized in Alg. 2, object candidates are detected on RGB frames by YOLO and temporally associated with ByteTrack to yield stable 2D tracks. For 3D localization,

each track is refined by a lightweight center-seeded breadth-first search on the depth image, expanding a mask of depth-consistent pixels. The resulting region is back-projected through camera intrinsics and extrinsics to obtain a 3D point cloud, from which both the centroid and an effective radius are estimated.

Obstacle velocities are then computed via temporal differencing, completing the state estimate $(\hat{p}_C^O, \hat{v}_C^O, \hat{r}_C^O, id)$. This design enables the camera pathway to complement LiDAR by providing high-resolution 3D localization and geometric size estimates whenever the obstacle is in view, while remaining efficient enough for real-time onboard deployment.

Short-horizon prediction. To anticipate obstacle motion, we employ an LSTM predictor trained on simulated trajectories. The network takes as input a short history of obstacle positions and velocities $\{p_{(t-k:t)}^O, v_{(t-k:t)}^O\}$ over several frames, and outputs one-step forecasts $\{\hat{p}_{t+1}^O, \hat{v}_{t+1}^O\}$. These predictions align with the perception outputs and provide forward-looking estimates of obstacle states under the constrained reaction horizon.

C. Threat-aware avoidance policy

The controller uses the robot state s_t^R and the perceived obstacle state $\{\hat{p}_t^O, \hat{v}_t^O, \hat{r}_t^O, id\}$ together with the unified threat \mathcal{T}_t . A monotone schedule $\beta_t = g(\mathcal{T}_t) \in [0, 1]$ continuously blends a conservative navigation controller with an aggressive reflex policy, while a lightweight turning module keeps the camera aligned with the most threatening direction. Commands are finally blended as $a_t = (1 - \beta_t) a_t^{\text{nav}} + \beta_t a_t^{\text{reflex}}$.

Active reorientation policy. Active reorientation is tied to the threat gap between LiDAR’s omnidirectional estimate and the current fused threat, so it only acts when LiDAR reveals a more urgent direction. Let $\phi_t = \text{atan2}((\hat{p}_t^O - p_t^R)_y, (\hat{p}_t^O - p_t^R)_x)$ be the target direction and ψ_t^R the yaw in θ_t^R . We track the heading with a bounded yaw-rate command scaled by the gap:

$$\psi_t^{\text{ref}} = \text{sat}\left(k_\psi [\mathcal{T}_t^L - \mathcal{T}_t]_+ (\phi_t - \psi_t^R)\right), \quad (2)$$

where $[x]_+ = \max(x, 0)$ and $\text{sat}(\cdot)$ limits excessive spin and $k_\psi > 0$ is the yaw-rate gain. Thus, when the camera’s target already dominates ($\mathcal{T}_t \geq \mathcal{T}_t^L$), turning naturally fades; if LiDAR sees a higher threat elsewhere, the gap grows and the robot reorients.

Navigation-based avoidance policy. Navigation-based avoidance is a threat-aware reference-tracking behavior that enlarges clearance under moderate risk. Let $\hat{d}_t = (p_t^R - \hat{p}_t^O) / \|\hat{p}_t^O - p_t^R\|$ be the obstacle-to-robot direction. We set a planar velocity reference opposite to the approach,

$$v_{t,xy}^{\text{ref}} = -\left(k_r^{\min} + (k_r^{\max} - k_r^{\min}) g(\mathcal{T}_t^C)\right) \hat{d}_{t,xy}, \quad (3)$$

and track it in the locomotion stack. Here k_r^{\min} and k_r^{\max} are the lower and upper bounds of the velocity gain. As \mathcal{T}_t^C increases, the gain grows smoothly, yielding longer, faster steps; LiDAR-driven updates ensure that the commanded velocity adaptively reorients according to the threat direction.

Reflexive evasion policy. When \mathcal{T}_t is high or the remaining reaction time is short, control blends toward a learned reflex




Fig. 3: Simulation experiment results of reflexive policy showing different evasion gaits. ● static obstacle; ● dynamic obstacle. (a) Sidestep to avoid the obstacle. (b)–(c) Jump to avoid the obstacle. (d) Crouch down to avoid the obstacle.

policy, adapted from the REBot baseline [12], that produces rapid evasive maneuvers (e.g., sidesteps, crouches, jumps). Perception is augmented with a short-horizon prediction \hat{p}_{t+1}^O to compensate sensing/actuation delays, and response intensity scales with \mathcal{T}_t via $g(\mathcal{T}_t)$. The reflex objective is a compact, threat-aware composite reward

$$r_t = \lambda_{\text{safe}} r_t^{\text{safe}} + \lambda_{\text{dir}} r_t^{\text{dir}} + \lambda_{\text{ene}} r_t^{\text{ene}} + \lambda_{\text{stab}} r_t^{\text{stab}} + \lambda_{\text{rec}} r_t^{\text{rec}}. \quad (4)$$

Safety r_t^{safe} penalizes small predicted clearance to the obstacle, pushing the robot to maintain a safe margin even under short reaction time. Direction r_t^{dir} discourages velocity components toward the obstacle and encourages net displacement away from the approach direction, reducing circling or oscillatory motion. Energy r_t^{ene} suppresses unnecessary actuation to keep evasive responses efficient rather than brute-force. Stability r_t^{stab} maintains dynamic steadiness of the base and reliable support during rapid maneuvers, improving controllability and preventing tumbles. Recovery r_t^{rec} encourages a quick return to stable posture and low velocities once the threat subsides, enabling smooth handover back to normal operation.

All terms are uniformly scaled by $g(\mathcal{T}_t)$, while the re-orientation and navigation-based controllers remain purely reference-tracking and independent of this reward.

V. EXPERIMENT

We evaluate APREBot on a quadruped platform with fully onboard sensing and control. The study proceeds from experimental setup and tasks descriptions (Sec. V-A), through policy training in simulation and deployment on real robot (Sec. V-B), followed by main results analysis (Sec. V-C), and ablations (Sec. V-D) that isolate critical components.

Our assessment is organized around four research questions: **Q1**: Does LiDAR-based detection and tracking outperform purely reactive use of raw point-cloud scans? **Q2**: Does hierarchical perception—LiDAR for omnidirectional scanning

combined with camera-based active focusing—yield superior performance compared to LiDAR-only sensing? **Q3**: Does active reorientation toward the most threatening target improve performance, particularly for rear or blind-spot approaches? **Q4**: Given active reorientation, does threat-aware scheduling provide advantages over fixed rule-based triggering schemes (e.g., TTC thresholds)?

A. Experiment settings

Baselines. We compare against three representative LiDAR-only baselines, covering both geometry-driven and learning-based approaches: 1) *Raycast* [24]: a method that directly uses raw LiDAR point clouds to generate control commands, without detection or tracking. 2) *FAPP* [13]: a method that applies DBSCAN clustering and Kalman filtering for explicit obstacle detection and tracking. 3) *Omni* [1]: a learning-based method that directly processes spatio-temporal point clouds with the PD-RiskNet perception module, enabling end-to-end omnidirectional avoidance.

Tasks. Obstacles are introduced from diverse directions around the robot, including rear and lateral blind zones outside the forward camera’s field of view. Their motion patterns span both slow, steady approaches and sudden accelerations. The obstacle set includes humans, balls, and sticks. In each trial, initial position, approach direction, and velocity are randomized within predefined ranges to ensure a broad distribution of threat conditions.

Metrics. We evaluate performance on 4 system-level metrics:

- **Avoidance Success Rate (ASR).** Defined by $ASR = N_{\text{avoid}}/N_{\text{total}}$, where N_{avoid} is the number of successful trials (no collision and stable recovery) out of N_{total} total trials.
- **Minimum clearance d_{min} .** Computed as $d_{\text{min}} = \min_t (\|p_o(t) - p_r(t)\| - R_s)$, where $p_o(t)$ and $p_r(t)$ denote obstacle and robot base positions, and R_s is the combined safety radius.

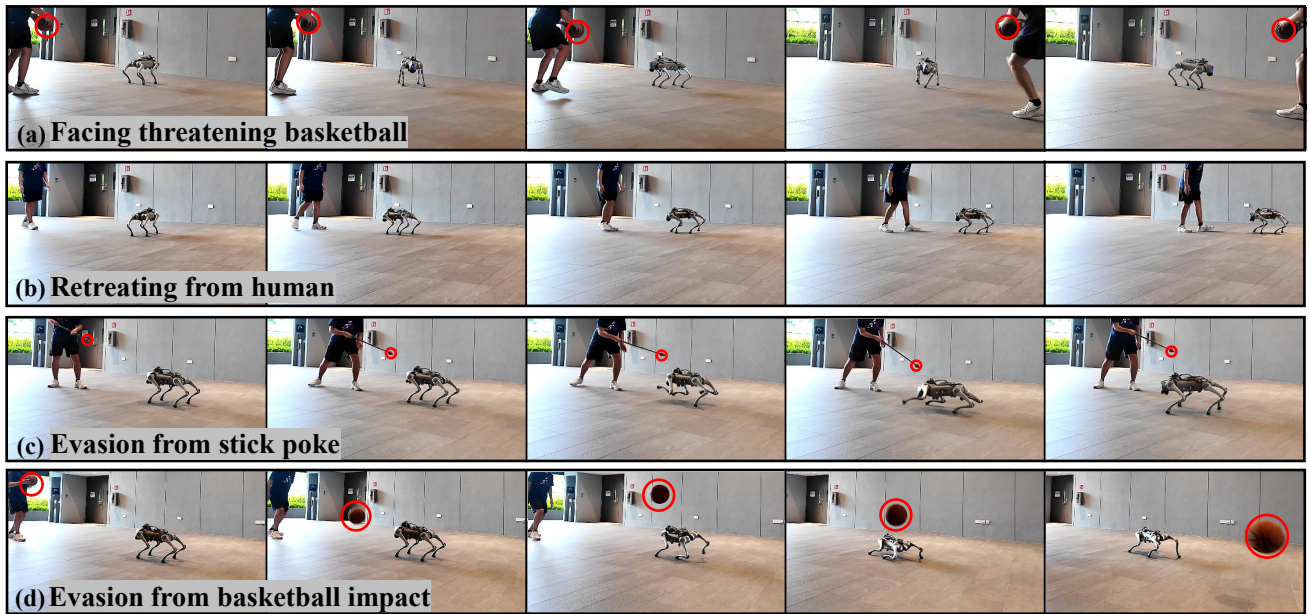


Fig. 4: Real-robot experiments demonstrating threat-aware decision making of the robot. (a) The robot identifies the basketball in the human’s hand as a threat and continuously reorients to face it. (b) The robot retreats slowly as the human approaches. (c) When poked rapidly with a stick, the robot performs reflexive evasion. (d) When hit with a basketball, the robot executes reflexive evasion.

- Normalized Trigger Lead (*TNL*). Defined by $TNL = (t_{\text{closest}} - t_{\text{trig}}) / \text{TTC}(t_{\text{trig}})$, where t_{trig} is the trigger time, t_{closest} the time of minimum distance, and $\text{TTC}(t_{\text{trig}})$ the estimated time-to-contact at trigger. We compute $\text{TTC}(t)$ under a constant-velocity assumption:

$$\text{TTC}(t) = \frac{\|\Delta p_t^s\|}{\max(0, v_{\text{rel}}^s(t) \cdot \hat{d}^s(t)) + \epsilon}, \quad (5)$$

where $\Delta p_t^s = p_t^O - p_t^R$ is the obstacle–robot displacement in the spatial frame, $\hat{d}^s(t) = \Delta p_t^s / \|\Delta p_t^s\|$ its unit direction, and $v_{\text{rel}}^s(t)$ the relative velocity. A small ϵ avoids division by zero.

- Energy E . The mechanical work between trigger and recovery: $E = \int_{t_{\text{trig}}}^{t_{\text{rec}}} |\tau(t)^\top \dot{q}(t)| dt \approx \sum_k |\tau_k^\top \dot{q}_k| \Delta t$, where τ and \dot{q} are joint torques and velocities.

B. Experiment Implementation

Policy training in simulation. We train the policy with PPO in Isaac Gym with 4,096 parallel environments using a Unitree Go2 quadruped model. Dynamic obstacles are modeled as spheres with randomized positions, velocities, and approach directions, ensuring diverse encounter conditions. Domain randomization is applied to robot dynamics (mass, friction, joint damping) and sensor noise, improving robustness for Sim2Real transfer. Policies converge stably within 6,000 iterations before deployment to the real robot.

The results show robust evasive maneuvers under diverse obstacle approaches (see Fig. 3). Depending on the incoming direction and remaining reaction time, the quadruped adapts with distinct strategies: lateral sidesteps (Fig. 3(a)), sudden jumps (Fig. 3(b)(c)), or crouching down (Fig. 3(d)) to avoid obstacles.

Real-robot deployment. Experiments are conducted on a Unitree Go2 quadruped equipped with onboard sensing and computation. The Livox MID-360 LiDAR provides 360° range data for obstacle detection and tracking, while the Intel RealSense D435 RGB-D camera offers high-resolution obstacle details. An onboard trajectory predictor estimates short-horizon obstacle trajectories, and all modules run in real time on the NVIDIA Jetson Orin without external sensing or offboard computation. Policies trained in simulation are deployed to the real robot directly.

Fig. 6 illustrates APREBot’s perception performance. LiDAR enables omnidirectional clustering and tracking of moving obstacles (Fig. 6(a)(b)), while the RGB-D camera segments objects such as humans and balls (Fig. 6(c)), and uses depth estimation to recover their 3D positions for precise localization (Fig. 6(d)). Fig. 4 shows representative trials on the quadruped platform. After LiDAR scanning, the robot actively reorients toward threatening obstacles (Fig. 4(a)). Under slow approaches it performs navigation-based retreat (Fig. 4(b)), whereas fast-approaching threats trigger reflexive evasions (Fig. 4(c)(d)), each followed by smooth recovery to its normal state.

C. Main results analysis

We evaluate three LiDAR-only baselines against APREBot on the real robot using the four metrics: ASR, d_{min} , TNL, and E . Tab. II summarizes the averaged outcomes across randomized trials, and Fig. 5 further illustrates performance differences. Several key findings emerge:

LiDAR detection improves over raw point cloud control. Fig. 5(a) shows that Raycast, which directly maps raw LiDAR scans to reactive actions, yields the smallest d_{min} and lowest ASR (Tab. II). By contrast, FAPP introduces explicit obstacle

TABLE II: Real-Robot Experiment Results

Method	ASR [%] \uparrow	d_{\min} [m] \uparrow	TNL \uparrow	E [J] \downarrow
Raycast	34.8 ± 10.2	0.095 ± 0.030	0.09 ± 0.05	1250 ± 220
FAPP	52.1 ± 9.8	0.140 ± 0.030	0.15 ± 0.07	890 ± 180
Omni	66.9 ± 8.9	0.135 ± 0.032	0.22 ± 0.06	930 ± 150
APREBot (ours)	78.7 ± 8.1	0.186 ± 0.035	0.30 ± 0.08	640 ± 140

ASR is computed over 50 trials, while d_{\min} , TNL, and E are averaged over successful avoidance trials only. Raycast: raw LiDAR point cloud control. FAPP: LiDAR clustering and tracking. Omni: LiDAR end-to-end RL. APREBot: our proposed method.

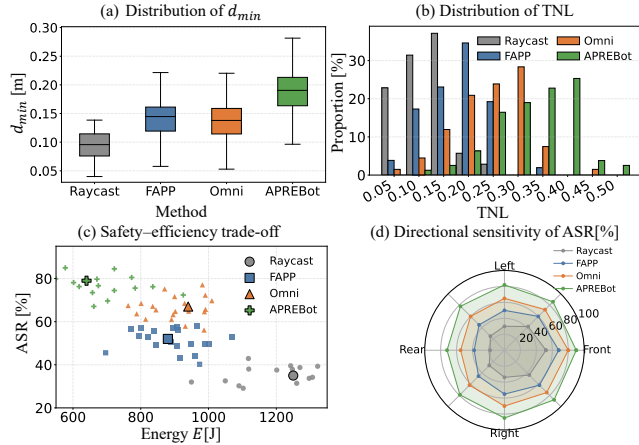


Fig. 5: Main experiment results. APREBot consistently outperforms LiDAR-only baselines with larger safety margins, earlier triggers, lower energy cost, and improved blind-side robustness.

clustering and tracking, significantly increasing ASR and enlarging d_{\min} . This confirms that structured LiDAR-based detection and tracking provides a more reliable safety margin than raw point cloud control.

Learning-based LiDAR methods improve success, while geometry-based detection is more efficient. As seen in Fig. 5(a) and Fig. 5(c), Omni achieves slightly higher ASR and TNL than FAPP, reflecting the strength of end-to-end LiDAR learning methods. However, FAPP maintains similar or better d_{\min} and lower energy consumption, indicating that geometry-based detection remains more efficient in terms of energy and stability (Fig. 5(b)(d)). The results reveal that learning and geometry-based approaches emphasize different priorities, underscoring the value of combining their strengths in future designs.

APREBot outperforms all LiDAR-only baselines. Compared to all baselines, APREBot consistently achieves the highest ASR and largest d_{\min} while reducing energy cost (Fig. 5(c), Tab. II). Fig. 5(b) shows that its TNL distribution shifts to the right, meaning evasive actions are triggered earlier and with larger safety margins. Moreover, Fig. 5(d) reveals that APREBot achieves superior ASR across all approach directions, with especially pronounced gains in rear and blind-side cases, highlighting the benefit of active reorientation and camera-based focusing.

D. Ablation studies

To assess the role of individual components, we compare APREBot against three ablated variants: removing trajectory




Fig. 6: Perception performance of APREBot. (a)–(b) LiDAR-based clustering and tracking of surrounding obstacles. (c) Camera RGB image with segmentation of a person and a basketball. (d) Camera depth image used for 3D position estimation.




Fig. 7: Ablation experiment results. APREBot achieves earlier triggering and larger safety margins compared to variants without prediction or threat, and maintains robust performance in rear and blind-side cases when active reorientation is removed.

prediction (*w/o Prediction*), disabling active reorientation (*w/o Reorient.*), and replacing threat-aware scheduling with a fixed TTC-based rule (*w/o Threat*). Tab. III reports averaged metrics, while Fig. 7 illustrates representative differences.

Prediction enables earlier responses. Fig. 7(a) shows that without prediction, the TNL distribution shifts leftward, indicating later activation. This leads to reduced ASR and smaller safety margins, as the robot occasionally fails against fast-approaching obstacles. With prediction, triggering occurs earlier and evasive maneuvers become more reliable, confirming its role in responsiveness.

Active reorientation ensures blind-side robustness. As seen in Fig. 7(b), disabling reorientation notably degrades

TABLE III: Ablation results on the experiments

Method	ASR [%] \uparrow	d_{\min} [m] \uparrow	TNL \uparrow	E [J] \downarrow
APREBot (ours)	78.7 ± 8.1	0.186 ± 0.035	0.30 ± 0.08	640 ± 140
w/o Prediction	74.5 ± 9.0	0.168 ± 0.045	0.23 ± 0.09	700 ± 150
w/o Reorient.	69.8 ± 9.6	0.148 ± 0.044	0.25 ± 0.08	780 ± 160
w/o Threat	72.9 ± 9.1	0.160 ± 0.050	0.22 ± 0.10	770 ± 170

w/o Prediction: removes short-horizon trajectory prediction. w/o Reorient.: disables threat-aware reorientation toward the most threatening target. w/o Threat: replaces threat-aware scheduling with a TTC-only triggering rule.

ASR in rear and lateral directions, while frontal performance remains closer to APREBot. By rotating toward the most threatening approach, APREBot brings obstacles into the camera field of view and maintains high success rates across all directions.

Threat-aware scheduling improves timing and efficiency.

Tab. III shows that replacing adaptive scheduling with a fixed TTC threshold lowers TNL and raises energy usage. As seen in Fig. 7(a), the fixed rule delays activation, so evasive actions occur later and often require more abrupt maneuvers. By contrast, threat-aware scheduling adapts to obstacle velocity and approach direction, enabling earlier and smoother responses.

In summary, prediction supports anticipation, reorientation provides blind-zone coverage, and threat-aware scheduling delivers well-timed and efficient reactions. Their combination gives APREBot the robustness that cannot be achieved by any single component alone.

VI. CONCLUSION

This work presented APREBot, a quadrupedal system that leverages active perception by integrating omnidirectional LiDAR scanning and camera-based active tracking to enable threat-aware decision making for reflexive evasion. By integrating real-time detection, short-horizon trajectory prediction, and threat-aware scheduling into the control policy, APREBot achieves agile and robust avoidance of dynamic obstacles without reliance on external devices. Extensive real-robot experiments demonstrated that APREBot substantially outperforms LiDAR-only baselines and validated its performance across diverse approach directions, including challenging rear and blind-side encounters. These results highlight the importance of active perception and adaptive decision making as foundations for robust quadrupedal reflexive evasion.

Future work will extend the APREBot framework to more general and demanding scenarios, such as evasion during locomotion and navigation in densely cluttered environments. These efforts aim to advance toward fully autonomous quadrupeds capable of safe and agile operation in unstructured and highly dynamic real-world settings.

REFERENCES

- [1] Z. Wang, T. Ma, Y. Jia *et al.*, “Omni-perception: Omnidirectional collision avoidance for legged locomotion in dynamic environments,” *arXiv preprint arXiv:2505.19214*, 2025.
- [2] W. Xu and F. Zhang, “Fast-lio: A fast, robust lidar-inertial odometry package by tightly-coupled iterated kalman filter,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3317–3324, 2021.
- [3] W. Xu, Y. Cai, D. He *et al.*, “Fast-lio2: Fast direct lidar-inertial odometry,” *IEEE Transactions on Robotics*, vol. 38, no. 4, pp. 2053–2073, 2022.
- [4] H. Wu, Y. Li, W. Xu *et al.*, “Moving event detection from lidar point streams,” *nature communications*, vol. 15, no. 1, p. 345, 2024.
- [5] T. He, C. Zhang, W. Xiao *et al.*, “Agile but safe: Learning collision-free high-speed legged locomotion,” *arXiv preprint arXiv:2401.17583*, 2024.
- [6] P. Roth, J. Nubert, F. Yang *et al.*, “Viplanner: Visual semantic imperative learning for local navigation,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 5243–5249.
- [7] S. Zhu, N. Perera, S. Yu *et al.*, “Dynamic object avoidance using event-data for a quadruped robot.”
- [8] S. Y. Alaba and J. E. Ball, “A survey on deep-learning-based lidar 3d object detection for autonomous driving,” *Sensors*, vol. 22, no. 24, p. 9577, 2022.
- [9] H. Liu, C. Wu, and H. Wang, “Real time object detection using lidar and camera fusion for autonomous driving,” *Scientific Reports*, vol. 13, no. 1, p. 8056, 2023.
- [10] B. Liu, C. Xu, W. Yang *et al.*, “Motion robust high-speed light-weighted object detection with event camera,” *IEEE Transactions on Instrumentation and Measurement*, vol. 72, pp. 1–13, 2023.
- [11] Z. Zhou, Z. Wu, R. Bouteau *et al.*, “Rgb-event fusion for moving object detection in autonomous driving,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 7808–7815.
- [12] Z. Xu, C. Hao, C. Wang *et al.*, “Rebot: Reflexive evasion robot for instantaneous dynamic obstacle avoidance,” *IEEE Robotics and Automation Letters*, vol. 11, no. 2, pp. 1746–1753, 2025.
- [13] M. Lu, X. Fan, H. Chen *et al.*, “Fapp: Fast and adaptive perception and planning for uavs in dynamic cluttered environments,” *IEEE Transactions on Robotics*, vol. 41, pp. 871–886, 2024.
- [14] X. Fan, M. Lu, B. Xu *et al.*, “Flying in highly dynamic environments with end-to-end learning approach,” *IEEE Robotics and Automation Letters*, 2025.
- [15] Y. Ma, A. Cramariuc, F. Farshidian *et al.*, “Learning coordinated badminton skills for legged manipulators,” *Science Robotics*, vol. 10, no. 102, p. eadu3922, 2025.
- [16] L. Schmid, O. Andersson, A. Sulser *et al.*, “Dynablox: Real-time detection of diverse dynamic objects in complex environments,” *IEEE Robotics and Automation Letters*, vol. 8, no. 10, pp. 6259–6266, 2023.
- [17] K. Koide, S. Oishi, M. Yokozuka *et al.*, “General, single-shot, targetless, and automatic lidar-camera extrinsic calibration toolbox,” *arXiv preprint arXiv:2302.05094*, 2023.
- [18] C. Yuan, X. Liu, X. Hong *et al.*, “Pixel-level extrinsic self calibration of high resolution lidar and camera in targetless environments,” *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 7517–7524, 2021.
- [19] H. Chen, Y. Lin, Z. Li *et al.*, “Irs: Instance-level 3d scene graphs via room prior guided lidar-camera fusion,” *arXiv preprint arXiv:2506.06804*, 2025.
- [20] Z. Zhuang, L. Zheng, W. Li *et al.*, “Rm-planner: Integrating reinforcement learning with whole-body model predictive control for mobile manipulation,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 7263–7269.
- [21] J. Lin, C. Zheng, W. Xu *et al.*, “R2live: A robust, real-time, lidar-inertial-visual tightly-coupled state estimator and mapping,” *arXiv preprint arXiv:2102.12400*, 2021.
- [22] J. Lin and F. Zhang, “R3live: A robust, real-time, rgb-colored, lidar-inertial-visual tightly-coupled state estimation and mapping package,” *arXiv preprint arXiv:2109.07982*, 2021.
- [23] Z. Xu, H. Shen, X. Han *et al.*, “Lv-dot: Lidar-visual dynamic obstacle detection and tracking for autonomous robot navigation,” *arXiv preprint arXiv:2502.20607*, 2025.
- [24] M. Pantic, I. Meijer, R. Bähnemann *et al.*, “Obstacle avoidance using raycasting and riemannian motion policies at khz rates for mavs,” *arXiv preprint arXiv:2301.08068*, 2023.