

HUNT: High-Speed UAV Navigation and Tracking in Unstructured Environments via Instantaneous Relative Frames

Alessandro Saviolo¹, Jeffrey Mao¹, and Giuseppe Loianno²

Abstract—Search and rescue operations require unmanned aerial vehicles to both traverse unknown unstructured environments at high speed and track targets once detected. Achieving both capabilities under degraded sensing and without global localization remains an open challenge. Recent works on relative navigation have shown robust tracking by anchoring planning and control to a visible detected object, but cannot address navigation when no target is in the field of view. We present HUNT (High-speed UAV Navigation and Tracking), a real-time framework that unifies traversal, acquisition, and tracking within a single relative formulation. HUNT defines navigation objectives directly from onboard instantaneous observables such as attitude, altitude, and velocity, enabling reactive high-speed flight during search. Once a target is detected, the same perception–control pipeline transitions seamlessly to tracking. Outdoor experiments in dense forests, container compounds, and search-and-rescue operations with vehicles and mannequins demonstrate robust autonomy where global methods fail. Video: <https://youtu.be/YsSf1qPPHhs>

I. INTRODUCTION

Unmanned Aerial Vehicles (UAVs), especially quadrotors, have become essential tools in Search-and-Rescue (SAR) operations, where their maneuverability and speed enable rapid deployment in unstructured, GPS-denied environments. These missions require two critical capabilities: the ability to *safely traverse unknown terrain at high speed* to search for victims or objects of interest, and the ability to *reliably track and follow detected targets* once they are identified. Achieving both in degraded sensing and unknown environments remains a fundamental challenge in robotics.

Traditional autonomy relies on global navigation anchored by GPS, Visual-Inertial Odometry (VIO), or Simultaneous Localization and Mapping (SLAM) [1]. These methods degrade in exactly the conditions most critical to SAR: dense forests, collapsed buildings, or urban canyons where GPS is degraded and visual features are sparse, dynamic, or occluded. Their dependence on persistent landmarks and loop closures makes global pose estimation fragile, resulting in drift and unsafe behavior at high speed.

An alternative paradigm, recently advanced in prior works [2], is to abandon global consistency and instead define navigation directly in a target-relative reference frame. By continuously re-anchoring perception, planning, and control

¹The authors are with New York University, NY 10012, USA. email: {as16054, jm7752}@nyu.edu.

²The author is with University of California Berkeley, Department of Electrical Engineering and Computer Sciences, Berkeley, CA 94720, USA. email: loiannog@eecs.berkeley.edu.

This work was supported by the NSF CAREER Award 2546659, the DARPA YFA Grant D22AP00156-00, and the DEVCOM ARL grant SARA W911NF-24-2-0057.



Fig. 1: HUNT enables reactive high-speed autonomy for search-and-rescue operations. The proposed framework is validated across diverse outdoor missions, from urban compounds to dense forests, demonstrating seamless integration of fast *traversal*, target *acquisition*, and sustained *tracking*.

to the instantaneous detection of a moving object, UAVs can robustly pursue targets even in unstructured or degraded environments. This *instantaneous relative navigation* formulation has demonstrated strong performance for the tracking phase of SAR missions. However, it presumes that a target is visible from the outset and remains continuously observable—an assumption that rarely holds in practice. What is missing is a capability for safe, high-speed traversal in unknown terrain *without any target directly in view*, and without reverting to the fragile assumptions of traditional GPS, VIO, or SLAM.

This paper introduces HUNT, a reactive autonomy framework that unifies high-speed traversal, acquisition, and tracking under a single instantaneous relative formulation. At its core is a novel *loitering mode*, which defines navigation objectives solely from directly observable quantities—attitude, altitude, and instantaneous velocity—allowing the UAV to traverse unknown cluttered environments safely even in the absence of a target. When a target is detected, the same perception–control pipeline seamlessly re-anchors the reference frame to the target, turning tracking into a special case of loitering without requiring global state.

Safety is enforced throughout by embedding high-order Control Barrier Functions (CBFs) directly into a Nonlinear Model Predictive Controller (NMPC), guaranteeing dynamically feasible, collision-free trajectories in real time. Robust transitions between modes are enabled by a confidence-based switching mechanism that prevents oscillations and ensures stability as detections appear and disappear.

We validate HUNT in extensive outdoor SAR-like missions spanning urban compounds, semi-structured layouts,

and dense forests. Experiments include loitering flights over city blocks and forest canopy, dense-clutter traversal below canopy, and full missions where the UAV searches, acquires, and pursues static and dynamic targets such as vehicles and mannequins. Across all settings, HUNT achieves consistent high-speed performance where global methods fail.

To the best of our knowledge, HUNT is the first framework to integrate traversal, acquisition, and tracking in cluttered, GPS-denied environments under a unified instantaneous relative formulation, eliminating reliance on global pose.

II. RELATED WORKS

Research on GPS-denied navigation can be broadly organized by the choice of reference frame. Early approaches adopted a *world-centric* formulation, estimating and controlling motion relative to a fixed global frame. VIO and Visual-Inertial SLAM (VI-SLAM) exemplify this class, fusing relative measurements over time to reconstruct trajectories [3]. While accurate in structured environments with abundant features and loop closures, they degrade rapidly at high speed or under feature sparsity. LiDAR integration has improved robustness [4], but such systems remain heavy, energy-inefficient, and fundamentally incompatible with agile flight. They are also susceptible to jamming [5] and interference in multi-agent scenarios [6]. Landmark-based localization [7] can reduce computation and support agile flight [8], yet presumes structured known environments unsuitable for SAR.

In contrast, *object-centric* formulations discard global consistency and express navigation relative to detected objects. A canonical example is Proportional Navigation [9], used since the 1950s for missile interception at extreme speeds. Variants have been deployed on quadrotors for agile pursuit [10] and adapted to camera-based control through image-based visual servoing for precision landing [11]. These methods relax environmental assumptions and support fast maneuvers, but are fragile to target loss when objects leave the field of view and cannot directly handle constraints.

Several efforts have sought to mitigate the loss of target visibility. Visibility-constrained methods maintain the target in the field of view [12] but require prior detection and cannot ensure observability without the object’s trajectory knowledge. Cooperative approaches leverage state broadcasts from tracked agents [13], but applicability is limited to collaborative scenarios. Maintaining a library of landmarks for relative navigation has been explored through online factor graphs [14] or predefined libraries [15], though these reintroduce world-centric challenges or demand prior maps. Open-loop fallback maneuvers have been demonstrated in structured contexts like competitive drone racing [16], and hybrid systems that switch between world- and object-centric formulations [17] improve overall system robustness but still suffer from the fragility of global pose estimation.

For handling constraints, optimization-based extensions have incorporated environmental constraints into relative navigation. Cooperative optimization has been demonstrated at low speeds with continuous state sharing [18], while

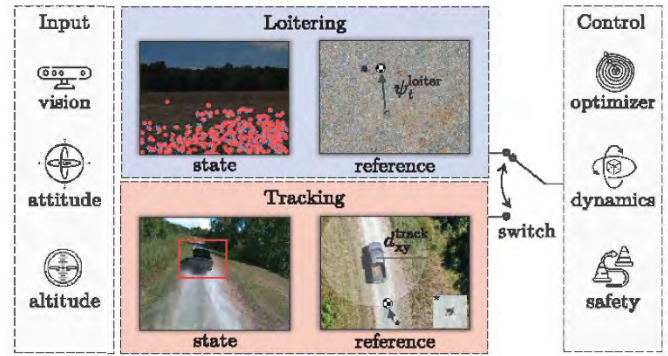


Fig. 2: **HUNT framework overview.** Inputs from onboard vision, attitude, and altitude sensors are fused to reconstruct directly observable states in the instantaneous relative frame \mathcal{R}_t . In *loitering*, \mathcal{R}_t is anchored to gravity and the initial heading, and the reference trajectory regulates altitude, heading (ψ_t^{loiter}), and forward velocity. In *tracking*, \mathcal{R}_t is re-anchored on the detected target, and the reference trajectory enforces a safety distance (d_{xy}^{track}) while continuously orienting toward it. A unified NMPC optimizes future trajectories under dynamics, actuation, and obstacle-avoidance constraints, with seamless switching between modes.

LiDAR-augmented schemes [19] expand awareness but assume static environments and modest velocities (< 1.5 m/s). Learning- and NMPC-based frameworks [2] address tracking in unstructured environments at higher speeds, yet all existing approaches remain fundamentally limited by their reliance on continuous target visibility.

III. METHODOLOGY

HUNT formulates perception, estimation, and control entirely within an *instantaneous relative frame* \mathcal{R}_t . This frame is reconstructed at every control cycle t from directly measurable quantities—attitude, altitude, instantaneous velocity, and relative target position when available—thereby eliminating dependence on unobservable states such as horizontal translation or absolute yaw. By anchoring objectives and constraints in \mathcal{R}_t , the system preserves the essential geometric structure required for constraint-aware control while avoiding the drift and fragility of global pose estimation.

Two operating regimes emerge naturally (Figure 2). In *loitering*, the UAV regulates its motion in an environment-anchored frame to traverse cluttered terrain while searching for a target. In *tracking*, the frame origin is re-anchored at each cycle to the detected target, yielding consistent pursuit. Both regimes are implemented within a single perception–control pipeline. The following subsections detail the construction of \mathcal{R}_t , the perception modules that supply instantaneous observables, and the control architecture that enforces dynamic feasibility and safety.

A. Loitering Mode

Loitering corresponds to the navigation regime where no target is detected with high confidence. In this setting, the UAV must sustain safe high-speed flight through cluttered

terrain while actively searching, without relying on global localization or precomputed maps. The instantaneous frame \mathcal{R}_t is anchored to gravity and the initial heading and then held fixed over time. Control in this regime depends only on states that are directly observable onboard and consistently mapped from the body frame \mathcal{B} into \mathcal{R}_t , ensuring robustness even as unobservable global variables drift.

1) *State Estimation*: State estimation is performed with an Unscented Kalman Filter (UKF) that maintains the state

$$\mathbf{x}_t = \begin{bmatrix} \mathbf{p}_t^{\mathcal{R}_t\mathcal{B}} \\ \mathbf{v}_t^{\mathcal{R}_t\mathcal{B}} \\ \mathbf{q}_t^{\mathcal{R}_t\mathcal{B}} \\ b_{p,t} \\ \mathbf{b}_{a,t} \\ \mathbf{b}_{g,t} \end{bmatrix}, \quad (1)$$

where $\mathbf{p}_t^{\mathcal{R}_t\mathcal{B}}$ is the relative position, $\mathbf{v}_t^{\mathcal{R}_t\mathcal{B}}$ the velocity, $\mathbf{q}_t^{\mathcal{R}_t\mathcal{B}}$ the attitude of \mathcal{B} with respect to \mathcal{R}_t , and $b_{p,t}$, $\mathbf{b}_{a,t}$, $\mathbf{b}_{g,t}$ the barometer, accelerometer, and gyroscope biases.

The state is propagated using IMU readings

$$\mathbf{z}_t^{\text{process}} = \begin{bmatrix} \mathbf{a}_t^{\mathcal{B}} \\ \boldsymbol{\omega}_t^{\mathcal{B}} \end{bmatrix}, \quad (2)$$

through inertial kinematics. Position $\mathbf{p}_t^{\mathcal{R}_t\mathcal{B}}$ is obtained by integrating velocity, $\mathbf{v}_t^{\mathcal{R}_t\mathcal{B}}$ by integrating accelerations $\mathbf{a}_t^{\mathcal{B}}$, and $\mathbf{q}_t^{\mathcal{R}_t\mathcal{B}}$ by integrating angular rates $\boldsymbol{\omega}_t^{\mathcal{B}}$.

The state is updated with the measurement vector

$$\mathbf{z}_t^{\text{update}} = \begin{bmatrix} p_{z,t}^{\mathcal{R}_t\mathcal{B}} \\ \mathbf{v}_t^{\mathcal{R}_t\mathcal{B}} \end{bmatrix}, \quad (3)$$

whose first component corresponds to the barometric altitude and second to the velocity reconstructed from stereo-inertial fusion. Specifically, velocity is obtained by tracking optical flow features between two sequential frames [20], rejecting outliers via essential-matrix RANSAC, and recovering metric scale from disparity. The resulting velocity is corrected for the camera lever arm, rotated into \mathcal{B} , and mapped into \mathcal{R}_t .

This choice of state follows the traditional formulation for UAV dynamics, allowing us to leverage decades of results in modeling, estimation, and control [21]. At the same time, it unavoidably introduces drift in the horizontal position components $p_{x,t}$, $p_{y,t}$, since they are not directly observable. In HUNT, this drift does not compromise navigation: reference generation (Subsection III-A.2) always defines goals relative to the current state estimate, ensuring that control objectives remain well-posed even as horizontal position diverges.

2) *Reference Generation*: Given the estimated state \mathbf{x}_t , loitering generates references in terms of altitude p_z^{loiter} , forward speed v^{loiter} , and heading ψ_t^{loiter} . The altitude reference regulates vertical motion, the forward speed sets the cruise velocity, and the heading defines the direction of travel.

The heading ψ_t^{loiter} maintains the initial yaw while continuously running the object detector. Low-confidence detections bias the heading toward the projected center of the bounding box, providing a weak directional cue without triggering a mode switch. Only when detection confidence

exceeds the switching threshold does the system transition to tracking, as detailed in Subsection III-E.

At each stage $i \in [0, N)$ of the NMPC prediction horizon of length N and discretization step Δt , the reference trajectory $\{\bar{\mathbf{x}}_{t+i}\}$ is rolled out by advancing at constant speed v^{loiter} along ψ_t^{loiter} while holding altitude at p_z^{loiter} . Formally,

$$\bar{\mathbf{p}}_{t+i} = \begin{bmatrix} p_{x,t} + (i v^{\text{loiter}} \Delta t) \cos \psi_t^{\text{loiter}} \\ p_{y,t} + (i v^{\text{loiter}} \Delta t) \sin \psi_t^{\text{loiter}} \\ p_z^{\text{loiter}} \end{bmatrix}, \quad (4)$$

$$\bar{\mathbf{q}}_{t+i} = \begin{bmatrix} \cos\left(\frac{\psi_t^{\text{loiter}}}{2}\right) \\ 0 \\ 0 \\ \sin\left(\frac{\psi_t^{\text{loiter}}}{2}\right) \end{bmatrix}, \quad (5)$$

corresponding to flight aligned with ψ_t^{loiter} . Because the horizontal position $p_{x,t}$, $p_{y,t}$ is included in the reference state, even if it drifts, the reference is defined relative to the current estimate, ensuring drift does not alter the commanded goal.

B. Tracking Mode

Tracking corresponds to the navigation regime where a target is detected with high confidence. In this setting, the UAV re-anchors the instantaneous frame \mathcal{R}_t at every cycle, placing its origin at the target position. This construction keeps the target stationary in \mathcal{R}_t regardless of global drift, ensuring that objectives are consistently defined in target-relative coordinates. Control in this mode aims to maintain a safety distance from the target, providing horizontal and vertical offsets while keeping the vehicle oriented toward it.

1) *State Estimation*: State estimation in tracking is performed with a UKF maintaining the same state vector as in loitering. The process model is driven by IMU angular rates

$$\mathbf{z}_t^{\text{process}} = [\boldsymbol{\omega}_t^{\mathcal{B}}], \quad (6)$$

since linear accelerations cannot be used to avoid placing any assumption about the knowledge of the target's motion. Attitude $\mathbf{q}_t^{\mathcal{R}_t\mathcal{B}}$ is therefore propagated from gyroscope readings, while position evolves by integrating velocity $\mathbf{v}_t^{\mathcal{R}_t\mathcal{B}}$.

The state is updated using target-relative displacement obtained from visual detection. A lightweight onboard detector with adaptive zoom runs on the RGB stream; crops tighten at high confidence and expand when uncertain to improve reacquisition (as detailed in prior works on target tracking [2]). Fused with stereo disparity, the detector output provides the target's 3D displacement in the camera frame, which is corrected for the camera lever arm, rotated into \mathcal{B} , and mapped into \mathcal{R}_t . The resulting measurement vector is

$$\mathbf{z}_t^{\text{update}} = [\mathbf{p}_t^{\mathcal{R}_t\mathcal{B}}], \quad (7)$$

where $\mathbf{p}_t^{\mathcal{R}_t\mathcal{B}}$ is the detected target position expressed in \mathcal{R}_t .

With this update, the UKF re-anchors the state to the target at every cycle. State and covariance are transported consistently into the moving frame, ensuring stable estimation.

2) *Reference Generation*: Given the estimated state \mathbf{x}_t , tracking generates references in terms of a horizontal stand-off distance d_{xy}^{track} , a vertical offset d_z^{track} , and a heading ψ_t^{track} that aligns the UAV with the target. The standoff regulates lateral separation, the vertical offset fixes relative altitude, and the heading ensures the target remains centered.

The heading ψ_t^{track} is updated at every cycle from the bearing of the detected target, reorienting the body x -axis to point toward it. This maintains a consistent look-at behavior.

At each stage $i \in [0, N)$ of the NMPC prediction horizon of length N and discretization step Δt , the reference trajectory $\{\bar{\mathbf{x}}_{t+i}\}$ is rolled out by holding the desired displacement and orientation fixed relative to the target. Formally,

$$\bar{\mathbf{p}}_{t+i} = \begin{bmatrix} p_{x,t}^{\mathcal{R}_t\mathcal{B}} \\ p_{y,t}^{\mathcal{R}_t\mathcal{B}} \\ 0 \end{bmatrix} \frac{d_{xy}^{\text{track}}}{\sqrt{(p_{x,t}^{\mathcal{R}_t\mathcal{B}})^2 + (p_{y,t}^{\mathcal{R}_t\mathcal{B}})^2}} + \begin{bmatrix} 0 \\ 0 \\ d_z^{\text{track}} \end{bmatrix}, \quad (8)$$

$$\bar{\mathbf{q}}_{t+i} = \begin{bmatrix} \cos\left(\frac{\psi_t^{\text{track}}}{2}\right) \\ 0 \\ 0 \\ \sin\left(\frac{\psi_t^{\text{track}}}{2}\right) \end{bmatrix}. \quad (9)$$

Since references are always expressed relative to the target in \mathcal{R}_t , horizontal drift has no effect.

C. Collision Avoidance

Collision avoidance is enforced by embedding CBF constraints directly in \mathcal{R}_t . At each cycle, stereo depth images are completed, projected into a point cloud [2], and the K most imminent collision points are selected using a time-to-collision heuristic [22]. Each sample k defines a local frame $\mathcal{O}_{k,t}$ and is first expressed in the camera frame \mathcal{C} as $\mathbf{o}_{k,t}^{\mathcal{C}\mathcal{O}_{k,t}}$. Using the calibrated extrinsics, this vector is mapped into the body frame \mathcal{B} and rotated into \mathcal{R}_t with $\mathbf{q}_t^{\mathcal{R}_t\mathcal{B}}$, yielding $\mathbf{o}_{k,t}^{\mathcal{R}_t\mathcal{O}_{k,t}}$.

For each candidate obstacle, the CBF encodes a quadratic separation constraint of the form

$$h_k(\mathbf{x}_t) = \|\mathbf{p}_t^{\mathcal{R}_t\mathcal{B}} - \mathbf{o}_{k,t}^{\mathcal{R}_t\mathcal{O}_{k,t}}\|^2 - r_{\text{safe}}^2, \quad (10)$$

where $\mathbf{p}_t^{\mathcal{R}_t\mathcal{B}}$ is the UAV position in \mathcal{R}_t and r_{safe} a safety radius. The admissible control set is defined by requiring

$$\ddot{h}_k(\mathbf{x}_t, \mathbf{u}_t) + 2\lambda\dot{h}_k(\mathbf{x}_t) + \lambda^2 h_k(\mathbf{x}_t) \geq 0, \quad (11)$$

for all k , ensuring forward invariance of the safe set and preventing collisions; $\lambda > 0$ is a tunable barrier rate, with larger values enforcing safety more aggressively and smaller values yielding smoother but less conservative behavior.

D. Motion Planning and Control

Both loitering and tracking are executed through a single unified controller: an NMPC augmented with CBF constraints. At each cycle, the NMPC solves a finite-horizon problem of length N , minimizing the deviation between the mode-specific predicted states $\hat{\mathbf{x}}_{t+i}$ and references $\bar{\mathbf{x}}_{t+i}$:

$$\min_{\substack{\hat{\mathbf{x}}_t, \dots, \hat{\mathbf{x}}_{t+N} \\ \hat{\mathbf{u}}_t, \dots, \hat{\mathbf{u}}_{t+N}}} \sum_{j=0}^N \|\hat{\mathbf{x}}_{t+j} - \bar{\mathbf{x}}_{t+j}\|_{\mathbf{Q}}^2 + \sum_{j=0}^{N-1} \|\hat{\mathbf{u}}_{t+j}\|_{\mathbf{R}}^2, \quad (12)$$

where \mathbf{Q}, \mathbf{R} are positive definite weighting matrices penalizing tracking error and control effort. The optimization is subject to dynamics, input bounds, and safety constraints for all predictions $j \in [0, N)$ and obstacles $k \in [0, K)$:

$$\hat{\mathbf{x}}_t = \mathbf{x}_t, \quad (13a)$$

$$\hat{\mathbf{x}}_{t+1+j} = f(\hat{\mathbf{x}}_{t+j}, \hat{\mathbf{u}}_{t+j}), \quad (13b)$$

$$\mathbf{x}_{\min} \leq \hat{\mathbf{x}}_{t+j} \leq \mathbf{x}_{\max}, \quad (13c)$$

$$\mathbf{u}_{\min} \leq \hat{\mathbf{u}}_{t+j} \leq \mathbf{u}_{\max}, \quad (13d)$$

$$\ddot{h}_k(\hat{\mathbf{x}}_{t+j}, \hat{\mathbf{u}}_{t+j}) + 2\lambda\dot{h}_k(\hat{\mathbf{x}}_{t+j}) + \lambda^2 h_k(\hat{\mathbf{x}}_{t+j}) \geq 0, \quad (13e)$$

where \mathbf{x}_t is the state estimate provided by the UKF without the biases and with angular velocity included, \mathbf{u}_t are individual motor thrusts, $f(\cdot)$ denotes the nominal dynamics [23], and the final inequality encodes the second-order CBF conditions (as derived in the collision avoidance Subsection III-C).

E. Mode Switching

Mode switching is handled through a hysteresis policy on the estimator uncertainty. At each cycle t , the tracking UKF provides a covariance matrix P_t over the target-relative state, and its trace $\text{tr}(P_t)$ is used as a scalar confidence measure:

$$\begin{aligned} \text{LOITER} \rightarrow \text{TRACK} : \quad & \text{tr}(P_t) \leq \tau_{\uparrow}, \\ \text{TRACK} \rightarrow \text{LOITER} : \quad & \text{tr}(P_t) \geq \tau_{\downarrow}, \end{aligned} \quad (14)$$

with $\tau_{\uparrow} < \tau_{\downarrow}$ enforcing hysteresis.

While in loitering, tentative detections can still bias the heading ψ_t^{loiter} toward the target without triggering a full switch. Only when estimator uncertainty falls below τ_{\uparrow} does the system re-anchor \mathcal{R}_t to the target and begin tracking. Conversely, if uncertainty exceeds τ_{\downarrow} , the system reverts to loitering. At each transition, the UKF and NMPC are reinitialized with the new frame definition and reference trajectory, ensuring consistent state propagation and planning. Because obstacle-avoidance constraints are recomputed from depth data at every cycle, CBF constraints remain valid.

IV. EXPERIMENTS

We evaluate HUNT in a range of real-world environments designed to stress its ability to loiter safely, track dynamic targets, and switch seamlessly between modes. Test sites include both structured settings (urban container compounds, open fields) and unstructured natural environments (dense forest canopy, narrow forest trails). Scenarios combine static clutter with dynamic elements such as off-road vehicles and human mannequins to emulate search-and-rescue conditions.

A. Experimental Setup

All experiments are conducted with a quadrotor weighing 1.2 kg, with a thrust-to-weight ratio of 4 and a motor span of 25 cm. Onboard computation is provided by an NVIDIA Jetson Orin NX, while low-level stabilization runs on a PixRacer Pro. The PixRacer's IMU and barometer, centrally mounted and coincident with the body frame \mathcal{B} , provide attitude, angular velocity, and altitude. Perception uses an Intel RealSense D455 operating at 320×240 fused with monocular disparity from DepthAnythingV2 [24] (ViT-S) accelerated

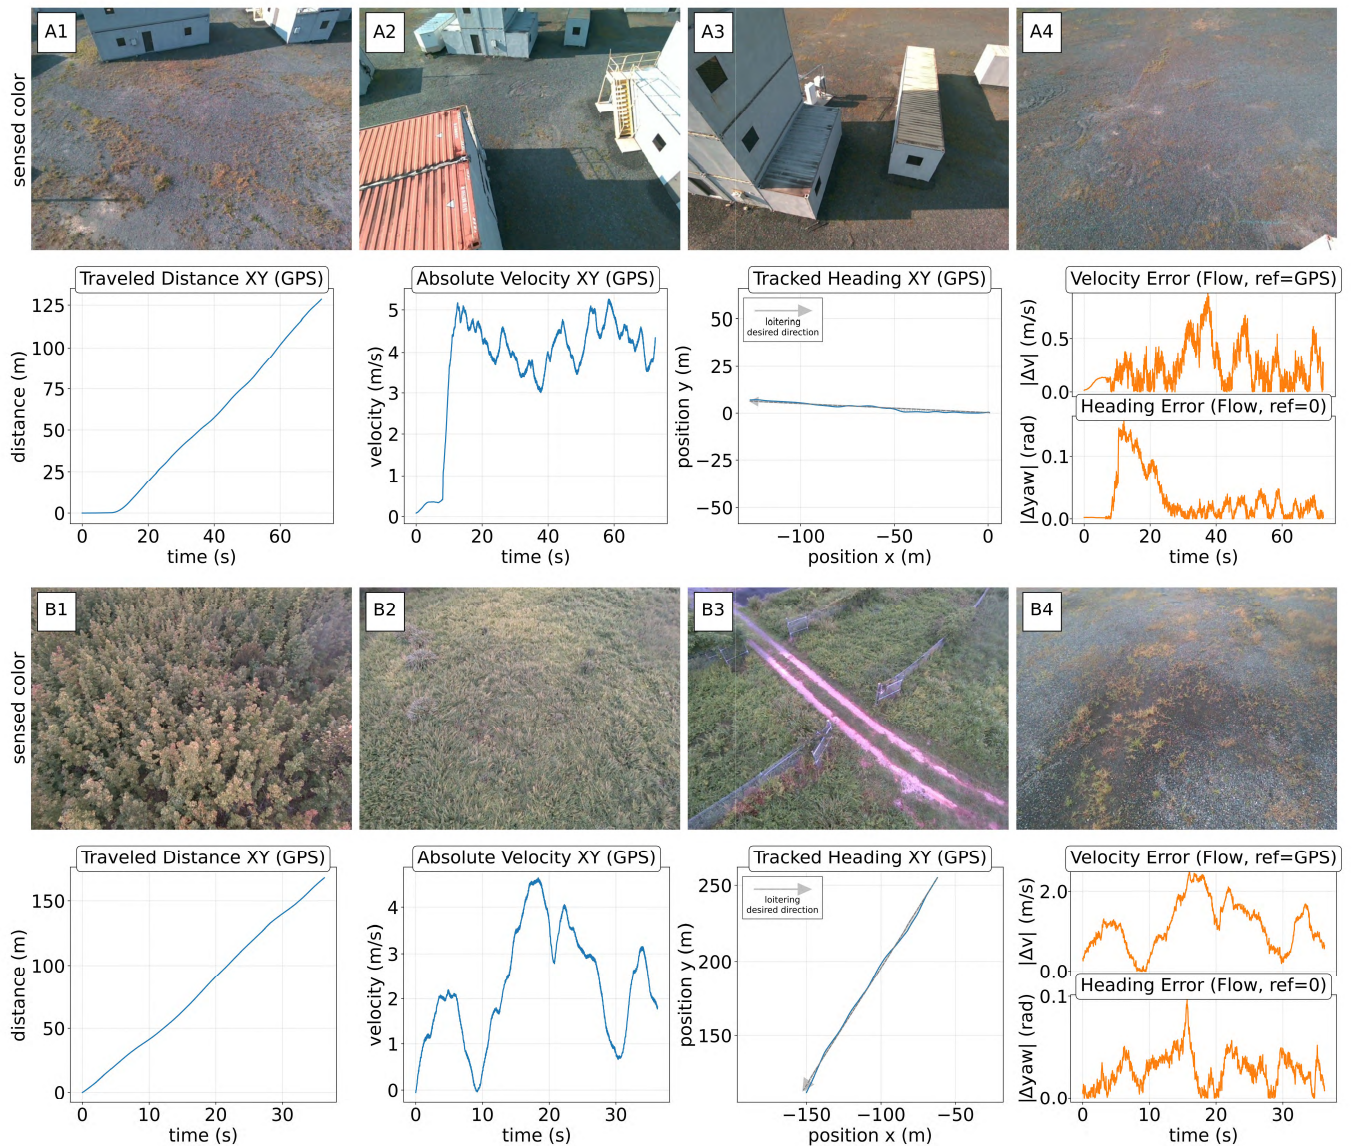


Fig. 3: **Loitering experiments without a target.** *Top:* Drift-free traversal above simulated city blocks. *Bottom:* Transition flight from forest canopy into the same city environment. In both experiments, UAV maintains steady velocity and heading using only instantaneous observables, with velocity and heading errors remaining bounded relative to GPS ground truth.

on TensorRT at 19 ms. Target detection is performed by a custom-trained YOLOv11 network, following the work [2]. The camera is mounted downward at 45° when flying above canopy or buildings, and forward-facing when operating at ground level to maximize the number of features in view.

The NMPC runs at 200 Hz with a 2 s prediction horizon and $N = 10$ shooting steps, solved with sequential quadratic programming in `acados` [25]. At each cycle, $K = 10$ high-risk obstacle points are selected based on minimal time-to-collision [22]. CBF constraints are updated at 60 Hz.

B. Loitering Above Structured and Natural Environments

We first evaluate HUNT in the loitering regime, where no target is present and the UAV must sustain safe forward flight while continuously searching (Figure 3). At mission start, the yaw reference is initialized to the vehicle’s takeoff heading ψ^{loiter} , the altitude reference is ramped to $p_z^{\text{loiter}} = 12$ m,

and forward motion begins at the commanded cruise speed $v^{\text{loiter}} = 4.0$ m/s. To maximize ground coverage, the RGB-D camera is pitched downward by 45° .

In the first experiment, the UAV flies above an emulated urban environment. As shown in the top row of Figure 3, the vehicle maintains steady forward progress at the commanded cruise speed, with measured velocity averaging 4.1 m/s and varying less than 0.3 m/s. Heading error relative to GPS ground truth remains bounded within 2° .

The second experiment repeats the mission in a more challenging mixed scenario, transitioning from flight above a dense forest canopy into structured city blocks. Despite the abrupt change in geometry and texture, the UAV preserves stable forward motion and heading traces (bottom row of Figure 3), showing that HUNT generalizes seamlessly across drastically different conditions.




Fig. 4: **Loitering under dense forest canopy.** *Left:* UAV trajectories for two representative runs, with GPS-reported position over time colored by estimated horizontal error (EPH). Despite errors in the order of 2–6 m that would make map-based navigation unreliable, HUNT sustains drift-free flight by relying only on instantaneous observables to guide it to maximum free space. *Right:* satellite view of the experimental site with reconstructed flight paths. Insets illustrate the camera views.

C. Safely Traversing Dense Unknown Forests

We next evaluate HUNT in loitering mode below a dense forest canopy, where clutter and degraded GPS make global localization unreliable (Figure 4). In this setting, the RGB-D camera is mounted forward-facing. The altitude reference is set to $p_z^{\text{loiter}} = 2.0$ m and forward speed to $v^{\text{loiter}} = 2.0$ m/s, while the yaw reference ψ_t^{loiter} is updated online at every cycle as the direction of maximum free space. Free space is computed by extracting the largest connected component in the completed depth map and selecting the point farthest from its boundary, yielding a heading that steers the UAV through the widest corridor.

Figure 4 presents two representative runs. The UAV sustains stable forward motion while weaving between trees, consistently maintaining clearance margins above 1.0 m. The reactive CBF-based avoidance ensures dead ends are bypassed and collisions averted even in dense clutter. By contrast, GPS logs—recorded only for evaluation—exhibit horizontal errors of 2–6 m and spurious altitude jumps. Such errors would prevent any GPS-based approach from constructing a reliable map and would likely lead to unsafe planning and eventual crashes.

D. Complete SAR Missions: Traversal, Acquisition, Tracking

We finally evaluate HUNT in three full SAR missions that require seamless transitions between loitering and tracking (Figure 5, Table I). In all cases, the UAV begins in loitering mode with fixed yaw reference ψ^{loiter} aligned to the initial heading at takeoff. Upon reliable detection, \mathcal{R}_t is re-anchored on the target and the system switches into tracking.

The first mission (*Urban Search*) emulates locating a stationary vehicle wreckage in an urban layout. The UAV loiters with altitude $p_z^{\text{loiter}} = 12$ m and forward speed $v^{\text{loiter}} = 5$ m/s until the car is detected. After switching, the displacement error remains bounded within ± 0.5 m, with plots confirming drift-free traversal before acquisition and consistent regulation thereafter.

The second mission (*Urban Pursuit*) mimics reconnaissance against a moving human target in a city. The UAV

TABLE I: **Quantitative performance across SAR missions.**

Each mission stresses a different capability: crossing a urban layout and acquiring a stationary vehicle wreckage (top row of Figure 5), transitioning from high-altitude search to dynamic human pursuit in a city environment (middle row), and traversing dense forest to locate and lock onto a target (bottom row). Metrics include UAV forward velocity (mean and peak), control effort (mean and peak linear acceleration), and pitch angle (mean and peak). HUNT maintains safe and agile navigation in all settings without parameter retuning.

Mission	Velocity (m/s)		Effort (m/s ²)		Pitch (°)	
	Mean	Max	Mean	Max	Mean	Max
Urban Search	8.85	12.8	10.31	22.35	2.00	15.30
Urban Pursuit	7.72	11.4	10.69	15.49	6.03	19.74
Forest Search	6.52*	9.50*	10.63	31.11	1.90	11.31

*GPS velocity under canopy is unreliable due to signal degradation.

surveys with the same loitering parameters ($p_z^{\text{loiter}} = 12$ m, $v^{\text{loiter}} = 5$ m/s, fixed ψ^{loiter}) until detecting a mannequin mounted on an all terrain vehicle (ATV). HUNT maintains pursuit at velocities up to 11.4 m/s, keeping lateral error confined to a 1.2 m corridor despite irregular gravel paths and dynamic motion. Notably, before the confidence threshold for switching is reached, low-confidence detections bias ψ_t^{loiter} toward the target. This is evident in the absolute GPS traces, where the trajectory gradually bends toward the target direction during loitering and then transitions smoothly into tracking once confidence surpasses the threshold. This heading bias mechanism effectively steers the robot toward potential targets even when the initial heading is misaligned.

The third mission (*Forest Search*) represents a rescue scenario in natural clutter. The UAV loiters below canopy with altitude $p_z^{\text{loiter}} = 2$ m and forward speed $v^{\text{loiter}} = 2$ m/s, following its initial heading while weaving through dense trees. Once a mannequin enters view, the system switches into tracking and sustains pursuit despite severe reflections and heading corrections up to 25°, while preserving safe

clearance enforced by CBF constraints. The same heading bias behavior observed in the urban pursuit mission is also present in this case, though less pronounced and more clearly visible in the attached multimedia material.

V. DISCUSSION AND FUTURE WORKS

HUNT lays the foundation for instantaneous relative navigation in unstructured environments, showing that both traversal and tracking can be expressed without global pose or mapping. A central design choice is the selection of the yaw reference ψ_t^{loiter} during loitering, which directly influences how effectively a target can be acquired while traversing. In our experiments, we demonstrated that a fixed initial heading is sufficient for safe forward flight, since yaw drift is minimal in the absence of global orientation. We further showed that this heading can be biased by low-confidence detections, effectively steering the UAV toward potential targets, or adapted online based on free-space analysis, yielding more exploratory but also erratic behavior.

Future work will extend these strategies with semantic reasoning from large vision–language models that can run at low frequency but provide high-level guidance. For instance, in an indoor search-and-rescue scenario after an earthquake, semantics such as recognizing windows or doors could guide loitering into unexplored regions, improving coverage without requiring explicit maps.

Another important future direction is the introduction of lightweight short-term memory. Since HUNT operates without global pose or persistent maps, the UAV may risk loitering in cycles. A memory of recently traversed landmarks would allow the system to bias navigation away from revisited areas and toward novel regions, improving efficiency without abandoning the instantaneous relative formulation.

VI. CONCLUSION

This work presented HUNT, a real-time framework for high-speed unmanned aerial navigation and tracking that eliminates reliance on global pose estimation. By formulating objectives and constraints in an instantaneous relative frame, HUNT unifies loitering, acquisition, and tracking within a single perception–control pipeline. Extensive field experiments across structured, semi-structured, and natural environments demonstrated fast traversal, reactive obstacle avoidance, and robust mode switching, achieving reliable autonomy where GPS is unreliable.

REFERENCES

- [1] J. Mao, R. C. Srinivas, S. Nogar, and G. Loianno, “Time-optimized safe navigation in unstructured environments through learning based depth estimation,” *IEEE Transactions on Field Robotics*, 2026.
- [2] A. Saviolo and G. Loianno, “Nova: Navigation via object-centric visual autonomy for high-speed target tracking in unstructured gps-denied environments,” *IEEE Transactions on Field Robotics*, vol. 2, pp. 826–846, 2025.
- [3] C. Campos, R. Elvira, J. J. Gomez, J. M. M. Montiel, and J. D. Tardos, “ORB-SLAM3: An accurate open-source library for visual, visual-inertial and multi-map SLAM,” *IEEE Transactions on Robotics*, vol. 37, no. 6, pp. 1874–1890, 2021.
- [4] W. Xu, Y. Cai, D. He, J. Lin, and F. Zhang, “Fast-lid2: Fast direct lidar-inertial odometry,” *IEEE Transactions on Robotics*, vol. 38, no. 4, pp. 2053–2073, 2022.
- [5] X. Hu, T. Liu, T. Shu, and D. Nguyen, “Spoofing detection for lidar in autonomous vehicles: A physical-layer approach,” *IEEE Internet of Things Journal*, vol. 11, no. 11, pp. 20 673–20 689, 2024.
- [6] A. L. Diehm, M. Hammer, M. Hebel, and M. Arens, “Mitigation of crosstalk effects in multi-lidar configurations,” in *Electro-Optical Remote Sensing XII*, vol. 10796. SPIE, 2018, pp. 13–24.
- [7] J. Xiao, D. Tortei, E. Roura, and G. Loianno, “Long-range uav thermal geo-localization with satellite imagery,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023, pp. 5820–5827.
- [8] S. Li, E. van der Horst, P. Duernay, C. De Wagter, and G. C. de Croon, “Visual model-predictive localization for computationally efficient autonomous racing of a 72-g drone,” *Journal of Field Robotics*, vol. 37, no. 4, pp. 667–692, 2020.
- [9] S. A. Murtaugh and H. E. Criel, “Fundamentals of proportional navigation,” *IEEE Spectrum*, vol. 3, no. 12, pp. 75–85, 1966.
- [10] A. Saviolo, P. Rao, V. Radhakrishnan, J. Xiao, and G. Loianno, “Unifying foundation models with quadrotor control for visual tracking beyond object categories,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 7389–7396.
- [11] J. Thomas, G. Loianno, K. Daniilidis, and V. Kumar, “Visual servoing of quadrotors for perching by hanging from cylindrical objects,” *IEEE Robotics and Automation Letters*, vol. 1, no. 1, pp. 57–64, 2015.
- [12] J. Mao, S. Nogar, C. M. Kroninger, and G. Loianno, “Robust active visual perching with quadrotors on inclined surfaces,” *IEEE Transactions on Robotics*, vol. 39, no. 3, pp. 1836–1852, 2023.
- [13] A. Ahmad, D. B. Licea, G. Silano, T. Báča, and M. Saska, “Pacnav: a collective navigation approach for uav swarms deprived of communication and external localization,” *Bioinspiration & Biomimetics*, vol. 17, no. 6, p. 066019, 2022.
- [14] G. Ellingson, K. Brink, and T. McLain, “Relative visual-inertial odometry for fixed-wing aircraft in gps-denied environments,” in *2018 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, 2018, pp. 786–792.
- [15] D. R. Maggio, C. Mario, B. Streetman, T. Steiner, and L. Carlone, “Vision-based terrain relative navigation on high altitude balloon and sub-orbital rocket,” in *AIAA SCITECH 2023 Forum*.
- [16] L. Scarciglia, A. Paolillo, and D. Palossi, “A map-free deep learning-based framework for gate-to-gate monocular visual navigation aboard miniaturized aerial vehicles,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2025, pp. 9935–9941.
- [17] T. Jantos, M. Scheiber, C. Brommer, E. Allak, S. Weiss, and J. Steinbrener, “Aivio: Closed-loop, object-relative navigation of uavs with ai-aided visual inertial odometry,” *IEEE Robotics and Automation Letters*, vol. 9, no. 12, pp. 10 764–10 771, 2024.
- [18] B. Zhang, X. Chen, Z. Li, G. Beltrame, C. Xu, F. Gao, and Y. Cao, “Coni-mpc: Cooperative non-inertial frame based model predictive control,” *IEEE Robotics and Automation Letters*, vol. 8, no. 12, pp. 8082–8089, 2023.
- [19] B. Zhang, X. Chen, Q. Chen, C. Xu, F. Gao, and Y. Cao, “Global-state-free obstacle avoidance for quadrotor control in air-ground co-operation,” *IEEE Robotics and Automation Letters*, 2025.
- [20] A. Alfarano, L. Maiano, L. Papa, and I. Amerini, “Estimating optical flow: A comprehensive review of the state of the art,” *Computer Vision and Image Understanding*, p. 104160, 2024.
- [21] A. Saviolo and G. Loianno, “Learning quadrotor dynamics for precise, safe, and agile flight control,” *Annual Reviews in Control*, vol. 55, pp. 45–60, 2023.
- [22] A. Saviolo, N. Picello, J. Mao, R. Verma, and G. Loianno, “Reactive collision avoidance for safe agile navigation,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2025, pp. 16 125–16 132.
- [23] A. Saviolo, G. Li, and G. Loianno, “Physics-inspired temporal learning of quadrotor dynamics for accurate model predictive trajectory tracking,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10 256–10 263, 2022.
- [24] L. Yang, B. Kang, Z. Huang, Z. Zhao, X. Xu, J. Feng, and H. Zhao, “Depth anything v2,” *Advances in Neural Information Processing Systems*, vol. 37, pp. 21 875–21 911, 2024.
- [25] R. Verschuere, G. Frison, D. Kouzoupis, J. Frey, N. van Duijkeren, A. Zanelli, B. Novoselnik, T. Albin, R. Quirynen, and M. Diehl, “acados: A modular open-source framework for fast embedded optimal control,” *Mathematical Programming Computation*, 2021.

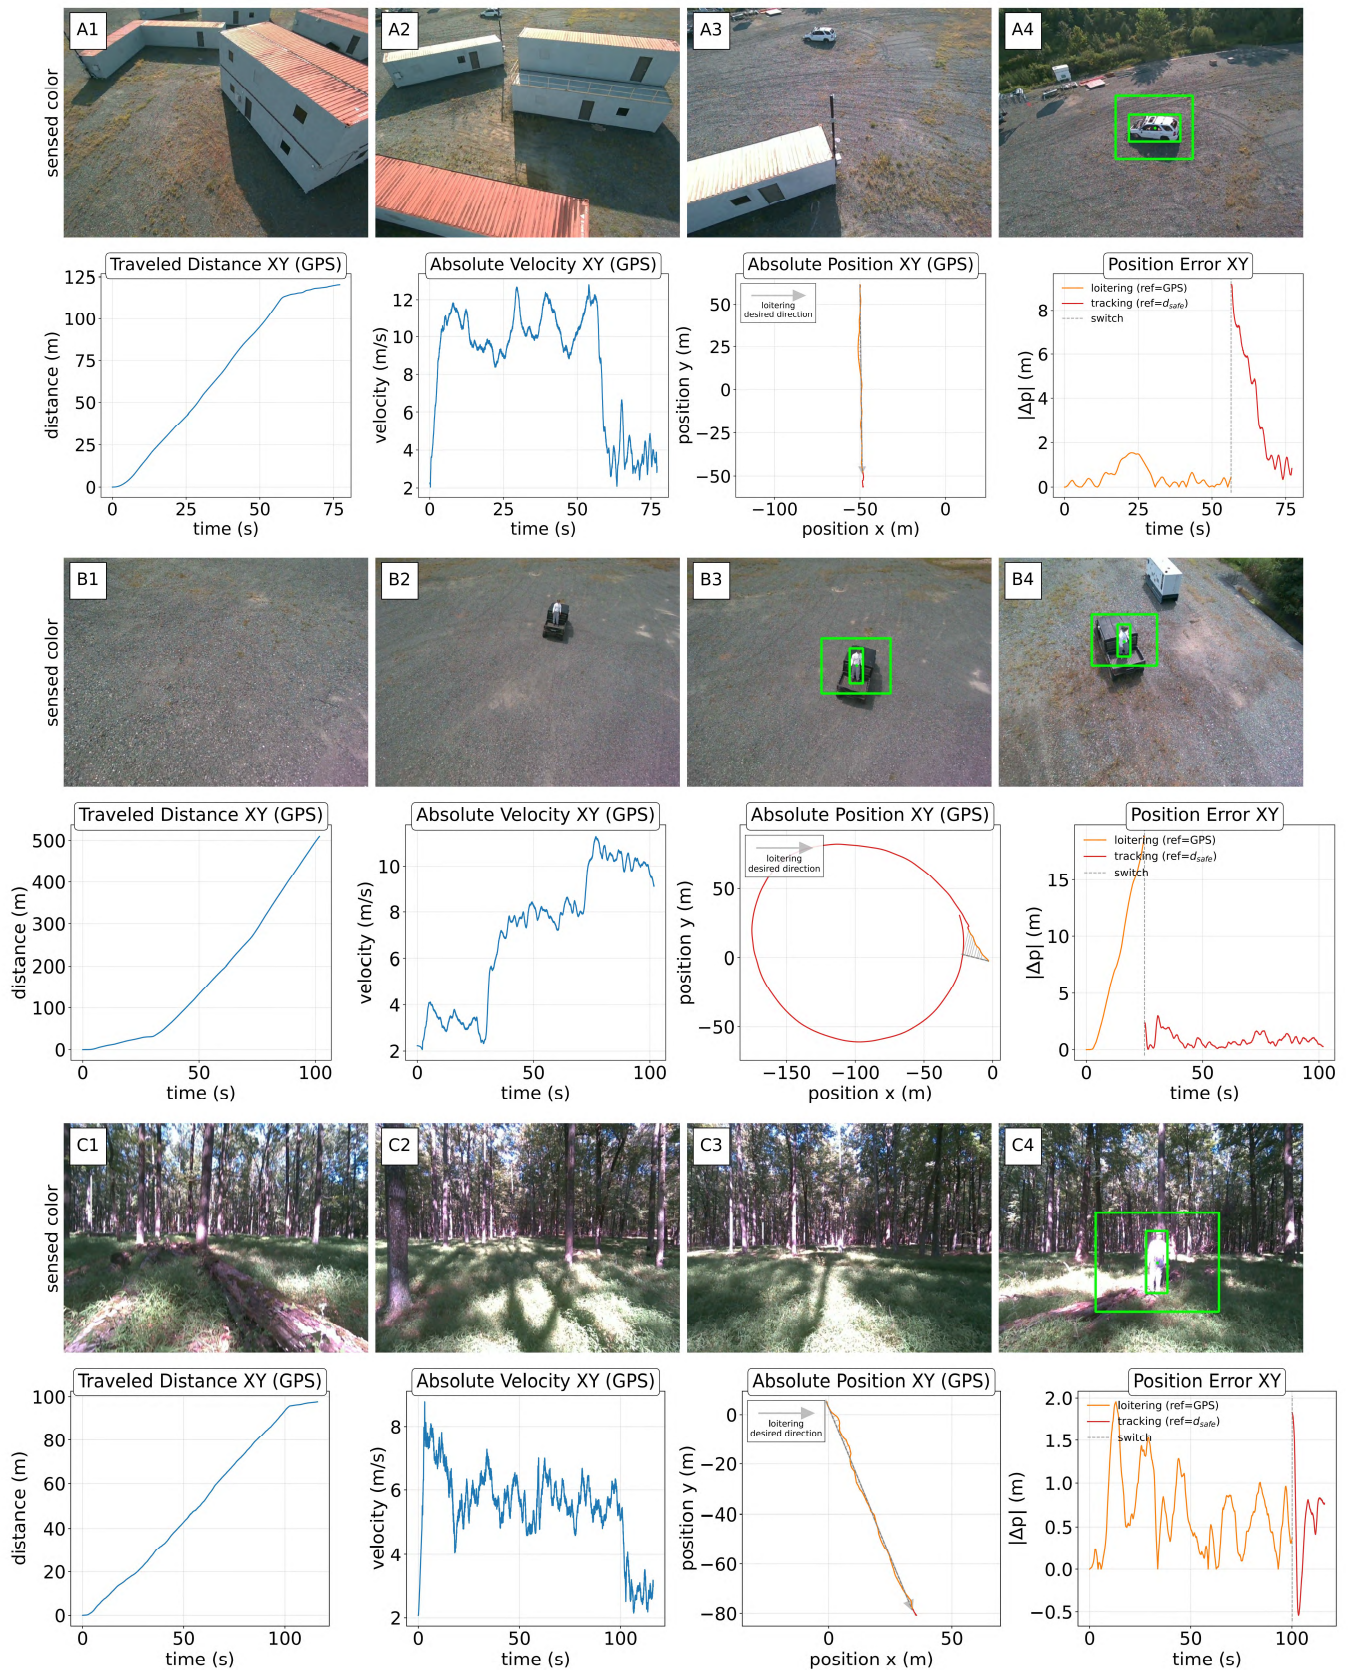


Fig. 5: Mode switching between loitering and tracking across three SAR missions. *Top (Urban Search)*: UAV flies above a simulated city, loiters while searching, then transitions into tracking upon detecting a stationary car wreckage. *Middle (Urban Pursuit)*: UAV surveys a city environment, acquires a mannequin on an ATV, and maintains pursuit along irregular gravel paths with stable error bounds. *Bottom (Forest Search)*: UAV loiters below dense forest canopy with active obstacle avoidance, then switches to tracking a mannequin despite clutter and occlusion.