

Eva-Tracker: ESDF-update-free, Visibility-aware Planning with Target Reacquisition for Robust Aerial Tracking

Yue Lin^{1,2}, Yang Liu^{1,3,*}, Dong Wang², Huchuan Lu^{2,3}

Abstract—The Euclidean Signed Distance Field (ESDF) is widely used in visibility evaluation to prevent occlusions and collisions during tracking. However, frequent ESDF updates introduce considerable computational overhead. To address this issue, we propose Eva-Tracker, a visibility-aware trajectory planning framework for aerial tracking that eliminates ESDF updates and incorporates a recovery-capable path generation method for target reacquisition. First, we design a target trajectory prediction method and a visibility-aware initial path generation algorithm that maintain an appropriate observation distance, avoid occlusions, and enable rapid replanning to reacquire the target when it is lost. Then, we propose the Field of View ESDF (FoV-ESDF), a precomputed ESDF tailored to the tracker’s field of view, enabling rapid visibility evaluation without requiring updates. Finally, we optimize the trajectory using differentiable FoV-ESDF-based objectives to ensure continuous visibility throughout the tracking process. Extensive simulations and real-world experiments demonstrate that our approach delivers more robust tracking results with lower computational effort than existing state-of-the-art methods.

I. INTRODUCTION

Maintaining visibility is essential for aerial autonomous tracking robots to continuously observe their targets. In cluttered environments, trackers must avoid collisions, prevent occlusions, maintain an appropriate tracking distance and observation angle, and promptly recover when the target is lost. Consequently, developing an efficient and robust tracking framework for aerial robots that ensures continuous visibility, handles target loss, and operates within limited onboard computational resources remains a significant challenge.

Numerous existing methods [1]–[6] prevent collisions by constraining the tracker within safe corridors. However, they often struggle to maintain target visibility in 3D environments, leading to frequent occlusions or suboptimal observation distances. The most widely adopted approaches [6]–[12] that jointly consider visibility and safety build on the Euclidean Signed Distance Field (ESDF). They leverage gradients of the ESDF to generate occlusion-free and collision-free trajectories. However, updating the ESDF is computationally intensive, which hinders real-time performance. Additionally, trajectory optimization typically covers only a small subregion of the ESDF update range, resulting in redundant computations [13]. More importantly, these methods frequently rely on complex multi-objective optimization, making it challenging to balance occlusion avoidance with maintaining an appropriate tracking distance.

¹Guangdong Laboratory of Artificial Intelligence and Digital Economy (SZ).
²School of Information and Communication Engineering, Dalian University of Technology. ³School of Future Technology, Dalian University of Technology. *Corresponding author: Yang Liu. Email: ly@dlut.edu.cn.




Fig. 1. Comparison between the traditional Euclidean Signed Distance Field (ESDF) and the proposed Field of View ESDF (FoV-ESDF). (a) The traditional ESDF encodes the distance from each voxel to its nearest obstacle across a wide spatial range. It must be updated in real time to reflect environmental changes, resulting in high computational overhead and substantial memory usage. (b) A slice of the FoV-ESDF, with white lines indicating the boundaries of the tracker’s FoV. The FoV-ESDF only encodes distances from voxels within the FoV to their nearest FoV boundary, while ignoring voxels outside the FoV. It is constructed once, requires significantly less memory, and eliminates the need for subsequent updates.

A promising approach to reducing the computational burden of ESDF-based methods is to utilize pre-built ESDFs, such as the RC-ESDF [14], which is specifically designed for collision assessment in robots with arbitrary shapes. By tailoring the ESDF to the robot’s geometry in advance, this method eliminates the need for updates. However, it primarily focuses on collision avoidance and does not incorporate visibility, which is crucial for aerial tracking tasks.

Inspired by the concept of pre-built ESDFs, we propose the Field of View ESDF (FoV-ESDF), a novel ESDF specifically designed for visibility in tracking scenarios, as illustrated in Fig. 1. Unlike the RC-ESDF, which focuses solely on collision avoidance, the FoV-ESDF explicitly models target visibility within the tracker’s FoV. It assigns gradients to obstacles located within the FoV, pushing them outward to preserve an unobstructed view, while ignoring those outside the FoV. Additionally, it generates a gradient for the target to maintain it centered and at an appropriate tracking distance, thereby ensuring continuous visibility. This visibility-aware design not only substantially reduces computational overhead but also unifies the optimization objectives of observation angle alignment, distance regulation, and occlusion avoidance through the FoV-ESDF, thereby significantly enhancing tracking robustness in cluttered 3D environments.

Initial path generation is another critical component of aerial tracking, as it provides an initial feasible solution for subsequent trajectory optimization. Many existing methods [1], [4], [10] rely on global searches, which incur high computational costs. Moreover, these methods fail immediately when the target is lost, lacking recovery mechanisms.

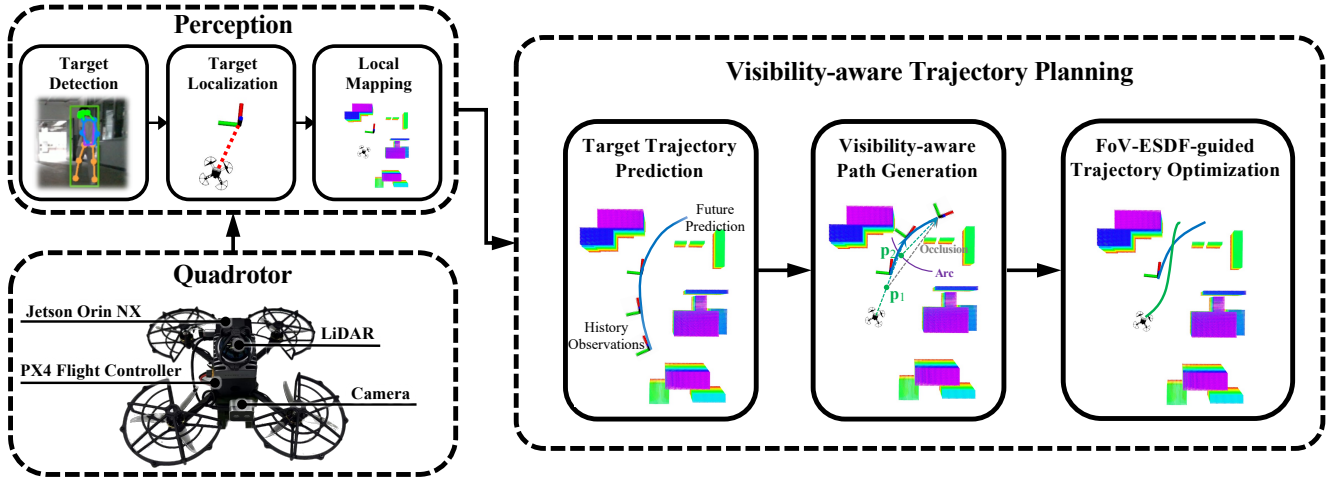


Fig. 2. Overview of the proposed system and the real-world drone platform. Left: The perception module processes sensor data for target detection, localization, and local mapping, all of which are executed entirely on the drone’s onboard computing device. Right: The trajectory planning module comprises three core components: target trajectory prediction, visibility-aware initial path generation, and FoV-ESDF-guided trajectory optimization. These components collaborate to produce occlusion-free tracking trajectories, ensuring continuous target visibility and enabling recovery when the target is lost.

To address these issues, we propose an efficient visibility-aware initial path generation algorithm that directly samples candidate observation points based on the predicted target position and the tracker’s pose, followed by visibility-aware adjustments to ensure an occlusion-free path. When the target is lost, the algorithm continues to generate paths using predicted positions, enabling prompt reacquisition. Compared to traditional methods, our algorithm significantly reduces computation time and enhances tracking robustness.

Building upon the FoV-ESDF and the initial path generation algorithm, we propose **Eva-Tracker**, an **ESDF**-update-free, **Visibility-Aware** trajectory planning framework for aerial tracking. As illustrated in Fig. 2, our system is composed of three tightly integrated components: target trajectory prediction, visibility-aware initial path generation, and FoV-ESDF-guided trajectory optimization. The prediction module generates smooth and continuous future target trajectories based on historical observations. The path generation module rapidly constructs occlusion-free initial paths that maintain an appropriate observation distance and support target reacquisition when the target is lost. Finally, the trajectory optimization module refines the tracking trajectory by maximizing visibility through a set of differentiable FoV-ESDF-based cost functions. Together, these modules enable robust and real-time aerial tracking in complex environments.

To evaluate the performance of our approach, we benchmark **Eva-Tracker** against several state-of-the-art tracking methods [1], [9], [10]. The experimental results demonstrate that **Eva-Tracker** achieves more robust tracking while significantly reducing computational costs. Furthermore, we deploy **Eva-Tracker** on a real-world quadrotor system and conduct tests in both indoor and outdoor environments. The results confirm that our method enables stable tracking even in challenging scenarios, highlighting its effectiveness and practicality in real-world applications.

In summary, the major contributions of this paper are:

- An efficient visibility-aware initial path generation algorithm with tracking recovery capability.
- A pre-built FoV-ESDF enables efficient visibility evaluation without requiring updates.
- A series of differentiable objective functions based on the FoV-ESDF to optimize the tracker’s pose.
- Comprehensive simulations and real-world experiments demonstrate superior tracking performance while reducing computational costs.

II. RELATED WORK

Recently, numerous learning-based trackers [15]–[20] have demonstrated impressive performance across various tracking tasks. However, they typically necessitate extensive training or environment-specific fine-tuning, and often generalize poorly to unfamiliar environments. Although **FAN** [21] improves robustness by integrating several large-scale models, its substantial computational overhead limits its applicability to real-time deployment on resource-constrained platforms.

To enhance adaptability in unknown environments and facilitate deployment on edge devices, a variety of modular methods have been explored. **Fast-Tracker** [2] combines a kinematic search front end with a trajectory optimization back end. However, it neglects visibility, often resulting in overly close tracking distances or target loss. **SF-Tracker** [10] maintains visibility by minimizing the observation angle error, but it does not account for occlusions. **Elastic Tracker** [1] models the visible area using 2D sectors, overlooking occlusions caused by 3D obstacles. In contrast, our method leverages the FoV-ESDF to simultaneously account for occlusions and observation distance in 3D environments, thereby ensuring continuous target visibility.

A widely adopted approach involves leveraging the ESDF for visibility optimization to avoid occlusions. The receding-horizon methods [7], [8] mitigate occlusion by maximizing the ESDF values along the line from the tracker to the target.

The visibility-aware tracking methods [9], [11] approximate the FoV using multiple spheres and maximize ESDF values of these spheres to prevent occlusion. However, constructing and updating the ESDF is time-consuming and computationally expensive. Moreover, these methods require additional objective functions to separately optimize tracking distance and observation angle. In contrast, our method employs a precomputed FoV-ESDF, eliminating the need for updates and enabling the joint optimization of occlusion, tracking distance, and observation angle within a unified framework.

Some studies have investigated target reacquisition. Fast-Tracker [2], [16] determines the navigation goal by predicting the target's future trajectory but does not consider visibility. ROAR [12] employs a Markov chain-based reacquisition mechanism. However, it estimates target motion solely from the current observation, which can result in significant prediction errors. In contrast, our method predicts the target's trajectory using a sequence of historical observations and explicitly incorporates visibility constraints.

III. VISIBILITY-AWARE PATH GENERATION

A. Target Trajectory Prediction

Given an observation time interval τ and a sequence of $m+1$ historical observation positions $\mathbf{h}_0, \mathbf{h}_1, \dots, \mathbf{h}_m \in \mathbb{R}^3$, we model the target's trajectory using a Bézier curve

$$\boldsymbol{\rho}(t) = \sum_{i=0}^n \frac{n!(T+t)^i(T-t)^{n-i}}{i!(n-i)!(2T)^n} \boldsymbol{\rho}_i, \quad t \in [-T, T], \quad (1)$$

where $n = 2m$, $T = m\tau$, and $\boldsymbol{\rho}_i \in \mathbb{R}^3$ represents the control point of the Bézier curve. The past observations correspond to $t \leq 0$, while future predictions correspond to $t > 0$.

To obtain the predicted trajectory, we solve the following constrained optimization problem:

$$\begin{aligned} \min_{\boldsymbol{\rho}} \quad & \int_{-T}^T \|\boldsymbol{\rho}^{(2)}(t)\|^2 dt, \\ \text{s.t.} \quad & \boldsymbol{\rho}(i\tau - T) = \mathbf{h}_i, \quad i \in \{0, 1, \dots, m\}. \end{aligned} \quad (2)$$

Smoothness is ensured by minimizing the high-order derivatives of the trajectory, while equality constraints enforce that the predicted trajectory passes through historical observation points. This problem can be reformulated as an equality-constrained convex quadratic programming (EQP) problem, which admits a closed-form solution [22].

B. Visibility-aware Initial Path Generation

To generate a suitable initial path for subsequent optimization, we propose a visibility-aware path generation algorithm that accounts for both observation distance and potential occlusions. This algorithm produces a sequence of waypoints $\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_m \in \mathbb{R}^3$, as illustrated in Fig. 3.

Given the predicted target position $\boldsymbol{\rho}(k\tau)$ and the tracker's previous position \mathbf{p}_{k-1} , we first enforce a preferred observation distance d . Specifically, we draw a ray from the target toward the tracker and select a candidate point \mathbf{c}_k along this ray at a distance d from the target:

$$\mathbf{c}_k = \boldsymbol{\rho}(k\tau) - d \frac{\mathbf{\Pi}_2(\boldsymbol{\rho}(k\tau) - \mathbf{p}_{k-1})}{\|\mathbf{\Pi}_2(\boldsymbol{\rho}(k\tau) - \mathbf{p}_{k-1})\|}, \quad (3)$$

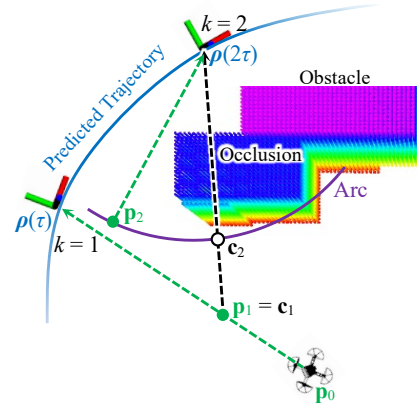


Fig. 3. Illustration of visibility-aware initial path generation.

where $\mathbf{\Pi}_2 = (\mathbf{e}_1, \mathbf{e}_2, \mathbf{0})$ is a projection matrix that extracts the horizontal component, and \mathbf{e}_i represents the i -th column of the third-order identity matrix.

Next, we check for occlusions along the line connecting \mathbf{c}_k and $\boldsymbol{\rho}(k\tau)$. As shown in Fig. 3 for $k=1$, if no occlusion is detected, \mathbf{c}_k is directly assigned as the waypoint \mathbf{p}_k . Conversely, as illustrated in Fig. 3 for $k=2$, if an occlusion is detected, we generate a circular arc around $\boldsymbol{\rho}(k\tau)$ with a radius d , represented as the purple arc in Fig. 3. Starting from \mathbf{c}_k , we perform a bidirectional search along the arc to identify the nearest occlusion-free point, which is then selected as \mathbf{p}_k . This process ensures that each waypoint maintains a clear line of sight to the target while preserving the desired observation distance, thereby providing a robust initial path for subsequent trajectory optimization.

In addition, this algorithm inherently enables target reacquisition. Even when the target is temporarily undetected, the system persists in generating occlusion-free waypoints at the optimal observation distance relative to the predicted target position. Once the target reappears within the FoV, the tracker seamlessly resumes normal tracking without the need for additional recovery mechanisms.

IV. FoV-ESDF-GUIDED TRAJECTORY OPTIMIZATION

A. FoV-ESDF

In this section, we detail the FoV-ESDF, which is precomputed offline based on the tracker's FoV, and does not require updates. Furthermore, it employs on-demand occlusion evaluation, concentrating solely on obstacles within the tracker's FoV, thereby significantly reducing computational overhead.

As illustrated by the black pyramid in Fig. 1(a), we model the FoV as a quadrilateral pyramid, denoted as

$$\left\{ (x, y, z) \mid 0 < x \leq D, \frac{|y|}{x} \leq \tan \frac{\alpha}{2}, \frac{|z|}{x} \leq \tan \frac{\beta}{2} \right\}, \quad (4)$$

where (x, y, z) denotes the coordinate in the tracker's frame, α and β represent the horizontal and vertical FoV angles of the camera, and D corresponds to the optimal observation distance d . Specifically, D is selected such that the point at $(d, 0, 0)$ achieves the maximum FoV-ESDF value.




Fig. 4. Illustration of FoV-ESDF-guided visibility optimization. Green circles represent obstacles within the FoV, while green arrows indicate the direction of the gradient descent generated by the FoV-ESDF for these obstacles. A red arrow indicates the direction of the gradient descent generated by the FoV-ESDF for the target. Orange circles represent obstacles that do not enter the FoV, which are excluded from the calculation.

For points within the FoV, the FoV-ESDF value is defined as the Euclidean distance to the nearest FoV boundary, while for points outside the FoV, it is set to zero. A planar cross-section of the FoV-ESDF is shown in Fig. 1(b). We construct the FoV-ESDF using the distance transform algorithm [23], and both its value and gradient at any point can be obtained through cubic linear interpolation [24]. Despite being pre-computed, the FoV-ESDF supports dynamic obstacle reasoning by transforming obstacle positions to the tracker's frame at each time step, enabling efficient visibility optimization.

To illustrate the optimization process of FoV-ESDF for visibility, we present an example in Fig. 4. Initially, some obstacles obstruct the target. The FoV-ESDF computes gradients for both the target and these obstacles, which are then propagated to the tracker's trajectory. The trajectory is subsequently optimized using gradient descent to maximize target visibility while minimizing occlusion. Specifically, the gradients generated for the obstacles are used to push them away from the FoV, while the gradients generated for the target are employed to maintain the target at the center of the FoV. After optimization, the obstacles are no longer within the FoV, and the target appears centered in the camera view.

B. Trajectory Representation and Problem Formulation

We parameterize the tracker's waypoints and yaw angles into a MINCO trajectory [25], formulated as

$$\mathfrak{T} = \{\mathbf{p}(t), \psi(t) \mid [\mathbf{C}, \Theta] = \mathcal{M}(\mathbf{P}, \Psi, \mathbf{t})\}, \quad (5)$$

where $\mathbf{p}(t) : \mathbb{R}_+ \mapsto \mathbb{R}^3$ and $\psi(t) : \mathbb{R}_+ \mapsto \mathbb{R}$ are polynomials of order $2s - 1$ with m segments, $\mathbf{P} = [\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_m] \in \mathbb{R}^{3 \times m}$ denotes the waypoints, $\Psi \in \mathbb{R}^{1 \times m}$ represents the yaw angles at these waypoints, the vector $\mathbf{t} = [t_1, t_2, \dots, t_m] \in \mathbb{R}_+^m$ specifies the duration of each trajectory segment, the tensor $\mathbf{C} = [\mathbf{C}_1, \mathbf{C}_2, \dots, \mathbf{C}_m] \in \mathbb{R}^{2s \times 3 \times m}$ and the matrix $\Theta = [\theta_1, \theta_2, \dots, \theta_m] \in \mathbb{R}^{2s \times m}$ contain the coefficients for the trajectory segments. The function \mathcal{M} maps \mathbf{P} , Ψ and \mathbf{t} to \mathbf{C} and Θ with linear complexity, allowing any second-order continuous cost function $J(\mathbf{C}, \Theta, \mathbf{t})$ to have a gradient applicable to the MINCO trajectory [25]. The i -th segment of \mathbf{p} and ψ is denoted as

$$\mathbf{p}_i(t) = \mathbf{C}_i^\top \beta(t), \quad \psi_i(t) = \theta_i^\top \beta(t), \quad (6)$$

where $\beta(t) = [1, t, t^2, \dots, t^{2s-1}]^\top$ represents the natural basis. The trajectory optimization problem is formulated as

$$\min_{\mathbf{P}, \Psi, \mathbf{t}} J_v(\mathbf{C}, \Theta, \mathbf{t}) + J_e(\mathbf{C}, \Theta, \mathbf{t}) + \lambda_p J_p(\mathbf{C}, \Theta, \mathbf{t}), \quad (7)$$

where J_v represents the visibility cost, J_e is the energy cost, J_p is the penalty term, with λ_p serving as its weighting factor. We solve this problem using L-BFGS algorithm [26] with Lewis-Overton line search [27].

C. Visibility Cost

1) *Overview*: The visibility cost is designed to ensure optimal observation by minimizing occlusion while maintaining a desirable observation angle and distance. It consists of three components, formulated as

$$J_v = \lambda_o \mathcal{J}_o + \lambda_v \mathcal{J}_v + \lambda_a \mathcal{J}_a, \quad (8)$$

where \mathcal{J}_o penalizes occlusions, \mathcal{J}_v enforces an optimal observation pose, and \mathcal{J}_a regulates the observation angle to facilitate convergence. The hyperparameters λ_o , λ_v , and λ_a control the relative importance of each term. Each component follows a unified numerical integration formulation:

$$\mathcal{J}_\mu = \sum_{i=1}^m \frac{t_i}{k_i} \sum_{j=0}^{k_i} \eta_j f_\mu \left(\mathbf{p}_i, \psi_i, \rho_i, \frac{j t_i}{k_i} \right), \quad \mu \in \{o, v, a\}, \quad (9)$$

where $(\eta_0, \eta_1, \eta_2, \dots, \eta_{k_i-1}, \eta_{k_i}) = (0.5, 1, 1, \dots, 1, 0.5)$, k_i represents the number of samples for the i -th numerical integration, and ρ_i denotes the target trajectory segment:

$$\rho_i(t) = \rho \left(\max \left\{ T, t + \sum_{k=1}^{i-1} t_k \right\} \right). \quad (10)$$

The gradients of the visibility cost are computed as

$$\frac{\partial \mathcal{J}_\mu}{\partial t_i} = \frac{1}{k_i} \sum_{j=0}^{k_i} \eta_j \left(f_\mu + \frac{j t_i}{k_i} \frac{\partial f_\mu}{\partial t} \right), \quad (11a)$$

$$\frac{\partial \mathcal{J}_\mu}{\partial \mathbf{C}_i} = \frac{t_i}{k_i} \sum_{j=0}^{k_i} \eta_j \frac{\partial f_\mu}{\partial \mathbf{C}_i}, \quad \frac{\partial \mathcal{J}_\mu}{\partial \theta_i} = \frac{t_i}{k_i} \sum_{j=0}^{k_i} \eta_j \frac{\partial f_\mu}{\partial \theta_i}. \quad (11b)$$

2) *Occlusion Penalty*: The occlusion penalty is designed to prevent obstacles from entering the FoV by penalizing their presence. It is defined as

$$f_o(\mathbf{p}_i, \psi_i, \rho_i, t) = \frac{1}{2} \left(\sum_{k=1}^{N_t} \Xi(\mathbf{v}_{k,t}) \right)^2, \quad (12a)$$

$$\mathbf{v}_{k,t} = \mathbf{R}(\psi_i(t))(\mathbf{w}_{k,t} - \mathbf{p}_i(t)), \quad (12b)$$

where $\Xi : \mathbb{R}^3 \mapsto \mathbb{R}$ represents the FoV-ESDF value, N_t is the number of obstacle points inside the FoV at time t , $\mathbf{w}_{k,t}$ is the k -th obstacle point in the world frame, $\mathbf{v}_{k,t}$ represents the coordinate of this point in the tracker's frame, and

$$\mathbf{R}(\theta) = \begin{bmatrix} \cos \theta & \sin \theta & 0 \\ -\sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (13)$$

Note that $\Xi(\mathbf{v}_{k,t})$ is precomputed offline, while the obstacle point $\mathbf{w}_{k,t}$ is updated online to handle dynamic occlusions.

Let $w = \sqrt{2f_o}$, and $\mathbf{d}_{k,t} = \mathbf{R}'(\psi_i(t))[\mathbf{w}_{k,t} - \mathbf{p}_i(t)]$, where $\mathbf{R}'(\theta)$ represents the derivative of the rotation matrix $\mathbf{R}(\theta)$. The gradients of the occlusion penalty f_o are computed as

$$\frac{\partial f_o}{\partial \boldsymbol{\theta}_i} = w \sum_{k=1}^{N_t} \beta(t) \mathbf{d}_{k,t}^\top \nabla \Xi(\mathbf{v}_{k,t}), \quad (14a)$$

$$\frac{\partial f_o}{\partial \mathbf{C}_i} = -w \sum_{k=1}^{N_t} \beta(t) \nabla^\top \Xi(\mathbf{v}_{k,t}) \mathbf{R}(\psi_i(t)), \quad (14b)$$

$$\frac{\partial f_o}{\partial t} = w \sum_{k=1}^{N_t} [\mathbf{d}_{k,t} - \mathbf{R}(\psi_i(t)) \mathbf{p}'_i(t)]^\top \nabla \Xi(\mathbf{v}_{k,t}), \quad (14c)$$

where $\nabla \Xi(\mathbf{v}_{k,t})$ represents the gradient of the FoV-ESDF at $\mathbf{v}_{k,t}$, which can be calculated by cubic linear interpolation.

3) *Observation Penalties*: To ensure optimal visibility, we optimize the tracker's observation pose by maximizing the target's FoV-ESDF value, formulated as

$$f_v(\mathbf{p}_i, \psi_i, \boldsymbol{\rho}_i, t) = \frac{1}{2} [\Xi(d\mathbf{e}_1) - \Xi(\mathbf{R}(\psi_i(t)) \boldsymbol{\delta}_i(t))]^2, \quad (15)$$

where $\boldsymbol{\delta}_i(t) = \boldsymbol{\rho}_i(t) - \mathbf{p}_i(t)$ represents the position of the target relative to the tracker, and $\Xi(d\mathbf{e}_1)$ is the maximum FoV-ESDF value. The gradient of this penalty can be computed in the same manner as f_o .

As shown in Fig. 1(b), due to the 3D nature of the FoV, the point at which the maximum value of the FoV-ESDF is obtained is not unique on a horizontal plane. To ensure that the target remains centered in the FoV, the optimal yaw angle of the tracker should be

$$\varphi_i(t) = \text{atan2}(\mathbf{e}_2^\top \boldsymbol{\delta}_i(t), \mathbf{e}_1^\top \boldsymbol{\delta}_i(t)). \quad (16)$$

To enforce this constraint, we design a straightforward angular penalty

$$f_a(\mathbf{p}_i, \psi_i, \boldsymbol{\rho}_i, t) = 1 - \cos \delta_i(t), \quad (17)$$

where $\delta_i(t) = \psi_i(t) - \varphi_i(t)$ represents the angle of the target relative to the tracker. The gradient of f_a is given by

$$\frac{\partial f_a}{\partial \boldsymbol{\theta}_i} = \beta(t) \sin \delta_i(t), \quad (18a)$$

$$\frac{\partial f_a}{\partial \mathbf{C}_i} = \frac{\sin \delta_i(t)}{\|\mathbf{\Pi}_2 \boldsymbol{\delta}_i(t)\|^2} \beta(t) \boldsymbol{\delta}_i^\top(t) (-\mathbf{e}_2, \mathbf{e}_1, \mathbf{0}), \quad (18b)$$

$$\frac{\partial f_a}{\partial t} = \sin \delta_i(t) \left[\psi'_i(t) - \frac{\mathbf{e}_3^\top (\boldsymbol{\delta}_i(t) \times \boldsymbol{\delta}'_i(t))}{\|\mathbf{\Pi}_2 \boldsymbol{\delta}_i(t)\|^2} \right]. \quad (18c)$$

D. Energy Cost and Penalty Functions

To maintain a smooth target view and minimize motion blur, we select $s = 3$, corresponding to jerk minimization [1]. Additionally, we incorporate a temporal regularization term to synchronize the duration of the tracker's trajectory with that of the target. Thus, the energy cost is formulated as $J_e =$

$$\sum_{i=1}^m \int_0^{t_i} \left[\|\mathbf{p}_i^{(s)}\|^2 + (\psi_i^{(s)})^2 \right] dt + \gamma \left(T - \sum_{i=1}^m t_i \right)^2, \quad (19)$$

where γ is the weight of the regularization term. The value and gradient of J_e can be computed analytically [25].

To ensure safety and kinematic feasibility, we introduce a numerical integration penalty term, formulated as

$$J_p = \sum_{i=1}^m \frac{t_i}{k_i} \sum_{j=0}^{k_i} \eta_j \sum_{\mu \in P} G_\mu \left(\mathbf{p}_i, \psi_i, \frac{j t_i}{k_i} \right), \quad (20a)$$

$$G_d(\mathbf{p}_i, \psi_i, t) = H(\mathbf{p}_i(t), \psi_i(t)), \quad (20b)$$

$$G_{p_h}(\mathbf{p}_i, \psi_i, t) = \max \left\{ \left\| \mathbf{p}_i^{(h)}(t) \right\|^2 - v_h^2, 0 \right\}, \quad (20c)$$

$$G_{a_h}(\mathbf{p}_i, \psi_i, t) = \max \left\{ \left(\psi_i^{(h)}(t) \right)^2 - \omega_h^2, 0 \right\}, \quad (20d)$$

where $h \in \{1, 2\}$, and $P = \{d, p_1, p_2, a_1, a_2\}$. G_d enforces safety constraints, the function H is a differentiable collision metric derived from RC-ESDF, with further details available in [14]. The hyperparameters v_1, v_2, ω_1 , and ω_2 represent the maximum velocity, acceleration, angular velocity, and angular acceleration of the tracker, respectively.

V. SIMULATIONS AND BENCHMARKING

A. Implementation Details

We benchmark Eva-Tracker against Elastic Tracker [1], Vis-Planner [9], and SF-Tracker [10]. These open-source methods represent state-of-the-art tracking performance. All trackers are configured with an identical FoV of $\alpha \times \beta = 69.4^\circ \times 42.5^\circ$, and the optimal observation distance is set to $d = 2.5\text{m}$. The maximum linear and angular velocities of the tracker are set to 2m/s and 1.5rad/s, respectively. During the experiment, the target's position is continuously transmitted to the tracker. The hyper-parameters of Eva-Tracker are set as follows: $\tau = 0.5\text{s}$, $m = 2$, $\gamma = \lambda_p = 10^5$, $\lambda_o = 100$, $\lambda_d = 10^3$, and $\lambda_a = 10^4$. All simulations are conducted on an Intel NUC Phantom Canyon equipped with an Intel Core i7-1165G7 CPU and an NVIDIA GeForce RTX 2060 GPU.

B. Benchmark Comparisons

As shown in Fig. 5, we construct a simulated environment with 100 randomly placed obstacles to evaluate the generality and robustness of the four methods. The target follows a 200-second trajectory involving sharp turns, movement toward the tracker, and traversal through cluttered regions.




Fig. 5. The random simulation environment and the target's trajectory.

TABLE I
TRACKING PERFORMANCE COMPARISON

Method	TD (m)	AE (rad)	OR(%)	FR(%)
SF-Tracker [10]	2.89 ± 0.59	0.14 ± 0.13	27.50	40.58
Vis-Planner [9]	2.87 ± 0.58	0.11 ± 0.08	17.32	28.50
Elastic tracker [1]	2.52 ± 0.41	0.09 ± 0.08	8.71	12.27
Eva-Tracker	2.50 ± 0.24	0.10 ± 0.08	4.36	6.75

We evaluate tracking performance of four methods using four metrics summarized in Tab. I: tracking distance (TD), yaw angle error (AE), occlusion rate (OR), and failure rate (FR). A tracking failure is defined as the target becomes occluded or leaves the tracker’s FoV. SF-Tracker exhibits frequent failures owing to its lack of optimization for occlusion and tracking distance. Vis-Planner experiences sub-optimal performance as a consequence of complex visibility-aware multi-objective optimization. Although Elastic Tracker effectively maintains accurate yaw through 2D sectors, it struggles with managing 3D occlusions. In contrast, Eva-Tracker consistently upholds the desired observation distance and achieves the lowest occlusion rate, thereby yielding the highest success rate among all evaluated methods.

To further analyze visibility, we present the spatial distribution of the target’s projection onto the tracker’s FoV in the x - y plane in Fig. 6. SF-Tracker lacks explicit optimization for tracking distance, exhibits a large standard deviation in tracking distance. Vis-Planner struggles to balance multiple visibility-aware objective functions, resulting in a scattered distribution. Both our method and Elastic Tracker maintain a compact and centered distribution. However, unlike Elastic Tracker, our approach explicitly accounts for 3D occlusions, thereby further enhancing visibility robustness.




Fig. 6. Comparison of the distribution heatmap of the target position relative to the tracker’s FoV on the x - y plane. The heatmaps illustrate the frequency of the target appearing at various locations, with a more concentrated distribution near the center indicating stable visibility. Warmer colors denote a higher frequency of the target at corresponding locations, and white triangles mark the boundaries of the tracker’s FoV.

TABLE II
COMPARISON OF AVERAGE COMPUTATION TIME

Method	Path (ms)	Optimize (ms)	Total (ms)
Elastic tracker [1]	0.56 ± 0.19	6.56 ± 0.92	7.12
Vis-Planner [9]	0.79 ± 0.78	35.94 ± 4.81	36.69
SF-Tracker [10]	0.93 ± 0.46	36.59 ± 5.60	37.52
Eva-Tracker	0.02 ± 0.01	10.73 ± 3.89	10.75

We evaluate the average computation time of these methods during tracking, as summarized in Tab. II. The proposed visibility-aware initial path generation algorithm is an order of magnitude faster than the others, producing feasible paths in significantly less time. Regarding trajectory optimization, Elastic Tracker achieves the fastest speed due to its simplified 2D-sector-based formulation, which omits 3D occlusion reasoning. In comparison, our Eva-Tracker fully accounts for visibility in 3D space and achieves the fastest optimization speed among ESDF-based approaches, such as Vis-Planner and SF-Tracker, as it does not require ESDF updates.

C. Visibility-aware Tracking in a 3D Scene

To effectively demonstrate Eva-Tracker’s capability to maintain visibility in intricate 3D environments, we design a representative scenario in which the target traverses an n-shaped tunnel and subsequently navigates around a two-layer 3D obstacle. The tracking trajectories of both Eva-Tracker and Elastic Tracker [1] are illustrated in Fig. 7.

The Elastic Tracker loses sight of the target within the tunnel, leading to tracking failure. In contrast, the Eva-Tracker utilizes FoV-ESDF to analyze 3D occlusions, ensuring continuous visibility throughout the process and demonstrating its superior ability to handle 3D structures.

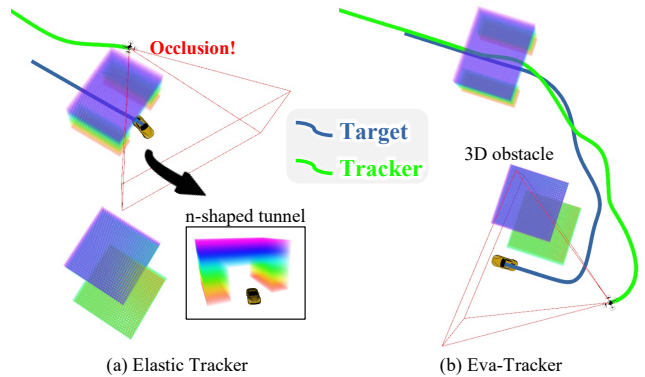


Fig. 7. Comparison of tracking trajectories in a 3D structured environment. (a) Elastic Tracker loses visibility when the target enters the n-shaped tunnel. (b) Eva-Tracker maintains clear visibility throughout, including while the target traverses the tunnel and navigates around a two-layer obstacle.

VI. REAL-WORLD EXPERIMENTS

To validate the practicality of Eva-Tracker, we deploy it on a real-world quadrotor platform, as shown in Fig. 2. The quadrotor is equipped with a Livox Mid-360 LiDAR for environmental perception and a camera for target detection. All computations are performed onboard using an NVIDIA Jetson Orin NX.

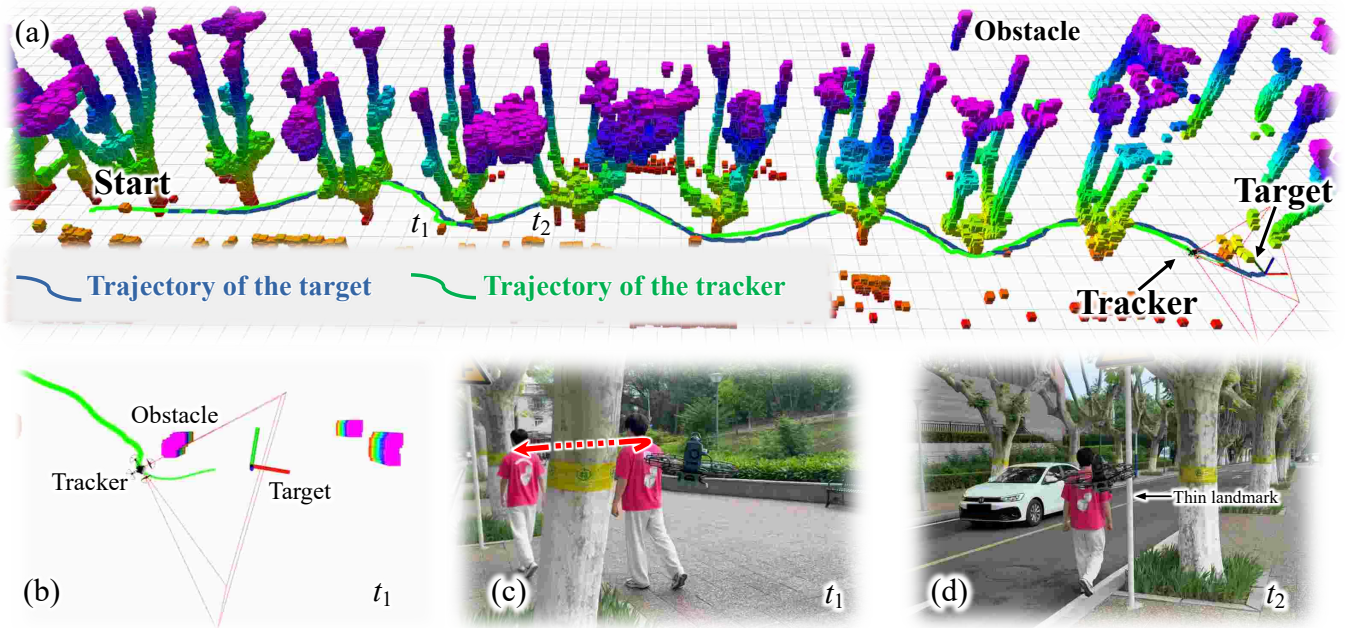


Fig. 8. Illustration of outdoor experiment. (a) Trajectories of the tracker and the target. (b) The planned trajectory of the tracker and the position of the target at time t_1 . (c) The tracker and the target at time t_1 . (d) The tracker and the target at time t_2 .

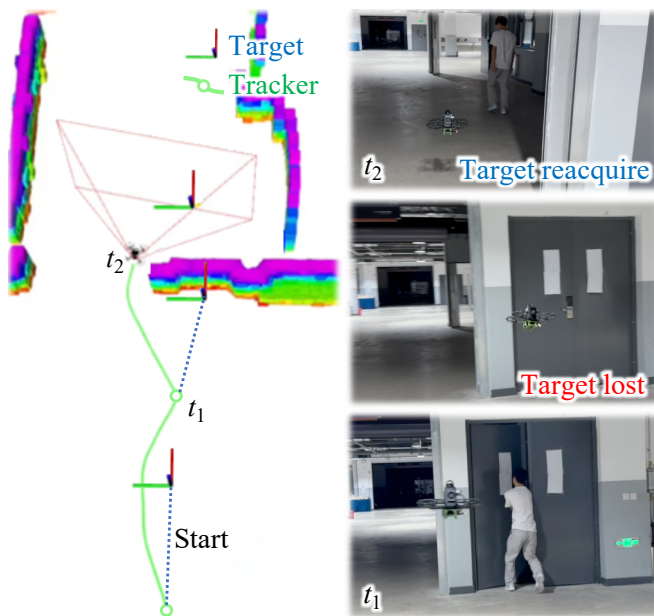


Fig. 9. Illustration of indoor experiment. The target walks through a door at time t_1 , after which the door closes and occludes the tracker’s view. The tracker proactively navigates around the obstacle, successfully reacquires the target at time t_2 , and resumes tracking.

We adopt Faster-LIO [28] for self-localization. The target is a person, whose pose is inferred by first detecting skeletal keypoints using YOLOv11 [29], followed by 3D pose estimation via MonoLoco [30]. When visual detection fails, the predicted position is used to attempt target reacquisition. To distinguish the target from obstacles, we remove LiDAR point clouds located within the estimated target radius prior to mapping and trajectory planning.

We evaluate Eva-Tracker in an unfamiliar outdoor environment populated with multiple trees. As shown in Fig. 8, the target repeatedly moves around tree trunks, which can easily cause occlusions. Throughout the experiment, the tracker proactively maneuvers around obstacles while maintaining an unobstructed line of sight to the target. The flight trajectory exhibits occlusion-free, collision-free, and smooth motion, with the tracker maintaining a stable tracking distance from the target. Notably, at time t_1 , the target suddenly changes direction after circumventing a large tree trunk. The tracker promptly adjusts its trajectory to avoid target loss. Additionally, at time t_2 , the tracker is able to adjust its trajectory to avoid occlusion even when the target passes by a thin landmark. This experiment demonstrates that Eva-Tracker can reliably maintain both visibility and safety during tracking.

To verify the robustness of Eva-Tracker in handling temporary target loss, we conduct an experiment in an unfamiliar, dynamic indoor environment, as shown in Fig. 9. The target walks through a narrow doorway, and the door subsequently closes, completely blocking the tracker’s line of sight. At time t_1 , the tracker uses the target’s estimated trajectory from historical observations to generate occlusion-free waypoints at the optimal observation distance relative to the predicted positions. This enables the quadrotor to navigate around the obstacle without hesitation. Upon reaching the other side of the doorway at time t_2 , the target reappears within the FoV, indicating successful tracking recovery, and the system seamlessly transitions back to normal tracking without requiring explicit recovery commands. The resulting trajectory confirms that Eva-Tracker can maintain operational continuity and rapidly reacquire the target in complex, dynamic environments where visibility is temporarily lost.

VII. CONCLUSION AND LIMITATIONS

In this paper, we propose Eva-Tracker, an ESDF-update-free, visibility-aware planning framework for aerial tracking that integrates a recovery-capable path generation mechanism. By leveraging a precomputed FoV-ESDF, Eva-Tracker eliminates the computational overhead associated with frequent ESDF updates. Combined with an efficient visibility-aware initial path generation algorithm, Eva-Tracker enables real-time performance and robust visibility maintenance in challenging 3D environments. Even during temporary target loss, Eva-Tracker facilitates rapid tracking recovery without requiring additional recovery procedures. Extensive simulations and real-world experiments demonstrate that Eva-Tracker outperforms state-of-the-art methods in both visibility preservation and recovery robustness in 3D scenarios.

Despite its advantages, Eva-Tracker has two primary limitations. First, the target reacquisition mechanism is effective only for short-term losses. If the target's actual trajectory deviates significantly from the predicted trajectory during prolonged occlusion, successful recovery becomes difficult. Additionally, the FoV-ESDF relies on a fixed camera FoV assumption, which limits its applicability to platforms equipped with dynamic FoV cameras such as zoom lenses or gimbal-mounted systems. Future work will address these limitations by developing a long-horizon motion prediction algorithm to manage extended target loss and by exploring adaptive FoV models to accommodate dynamic fields of view.

VIII. ACKNOWLEDGMENTS

This work is supported by the National Natural Science Foundation of China (Nos. 62293542, U23A20384, and 62476044), the Open Research Fund from Guangdong Laboratory of Artificial Intelligence and Digital Economy (SZ) (No. GML-KF-24-19), as well as the Open Research Fund from National Key Laboratory of China Space Intelligent Control Technology (No. HTKJ2024KL502016).

REFERENCES

- [1] J. Ji, N. Pan, C. Xu, and F. Gao, "Elastic tracker: A spatio-temporal trajectory planner for flexible aerial tracking," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2022, pp. 47–53.
- [2] Z. Han, R. Zhang, N. Pan, C. Xu, and F. Gao, "Fast-tracker: A robust aerial system for tracking agile target in cluttered environments," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 328–334.
- [3] Z. Zhang, Y. Zhong, J. Guo, Q. Wang, C. Xu, and F. Gao, "Auto filmer: Autonomous aerial videography under human interaction," *IEEE Robotics and Automation Letters*, pp. 784–791, 2023.
- [4] Q. Ren, H. Yu, Dai, Jiajun, Z. Zheng, J. Meng, L. Xu, C. Xu, F. Gao, and Y. Cao, "Intention-aware planner for robust and safe aerial tracking," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024, pp. 8153–8160.
- [5] Y. Lee, J. Park, B. Jeon, S. Jung, and H. J. Kim, "BPMP-tracker: A versatile aerial target tracker using bernstein polynomial motion primitives," *IEEE Robotics and Automation Letters*, pp. 10938–10945, 2024.
- [6] Q. Zou, H. Kong, and W. Yang, "REACTER: Perception-informed adaptive tracker in cluttered environments," *IEEE Robotics and Automation Letters*, pp. 9422–9429, 2025.
- [7] B. F. Jeon and H. J. Kim, "Online trajectory generation of a MAV for chasing a moving target in 3D dense environments," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 1115–1121.
- [8] B. Jeon, Y. Lee, and H. J. Kim, "Integrated motion planner for real-time aerial videography with a drone in a dense environment," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 1243–1249.
- [9] Q. Wang, Y. Gao, J. Ji, C. Xu, and F. Gao, "Visibility-aware trajectory optimization with application to aerial tracking," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 5249–5256.
- [10] Y. Lin, Y. Liu, P. Zhang, X. Chen, D. Wang, and H. Lu, "Safety-first tracker: A trajectory planning framework for omnidirectional robot tracking," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024, pp. 5416–5423.
- [11] Y. Gao, J. Ji, Q. Wang, R. Jin, Y. Lin, Z. Shang, Y. Cao, S. Shen, C. Xu, and F. Gao, "Adaptive tracking and perching for quadrotor in dynamic scenarios," *IEEE Transactions on Robotics*, pp. 499–519, 2024.
- [12] T. Zhang, C. Li, K. Zhao, H. Shen, and T. Pang, "ROAR: A robust autonomous aerial tracking system for challenging scenarios," *IEEE Robotics and Automation Letters*, pp. 7571–7578, 2025.
- [13] X. Zhou, Z. Wang, H. Ye, C. Xu, and F. Gao, "Ego-planner: An ESDF-free gradient-based local planner for quadrotors," *IEEE Robotics and Automation Letters*, pp. 478–485, 2020.
- [14] S. Geng, Q. Wang, L. Xie, C. Xu, Y. Cao, and F. Gao, "Robo-centric ESDF: A fast and accurate whole-body collision evaluation tool for any-shape robotic planning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023, pp. 290–297.
- [15] S. Bhagat and P. Sujit, "UAV target tracking in urban environments using deep reinforcement learning," in *International Conference on Unmanned Aircraft Systems (ICUAS)*, 2020, pp. 694–701.
- [16] N. Pan, R. Zhang, T. Yang, C. Cui, C. Xu, and F. Gao, "Fast-tracker 2.0: Improving autonomy of aerial tracking with active vision and human location regression," *IET Cyber-Systems and Robotics*, pp. 292–301, 2021.
- [17] H. Masnavi, J. Shrestha, M. Mishra, P. Sujit, K. Krusamäe, and A. K. Singh, "Visibility-aware navigation with batch projection augmented cross-entropy method over a learned occlusion cost," *IEEE Robotics and Automation Letters*, pp. 9366–9373, 2022.
- [18] P. Yan, J. Guo, X. Su, and C. Bai, "Long-term tracking of evasive urban target based on intention inference and deep reinforcement learning," *IEEE Transactions on Neural Networks and Learning Systems*, pp. 16 886–16 900, 2023.
- [19] H. Masnavi, A. K. Singh, and F. Janabi-Sharifi, "Differentiable-optimization based neural policy for occlusion-aware target tracking," *IEEE Robotics and Automation Letters*, pp. 11 714–11 721, 2024.
- [20] J. Lu, Y. Hui, X. Zhang, W. Feng, H. Shen, Z. Li, and B. Tian, "YOPOv2-tracker: An end-to-end agile tracking and navigation framework from perception to action," *arXiv preprint arXiv*, 2025.
- [21] A. Maalouf, N. Jadhav, K. M. Jatavallabhula, M. Chahine, D. M. Vogt, R. J. Wood, A. Torralba, and D. Rus, "Follow anything: Open-set detection, tracking, and following in real-time," *IEEE Robotics and Automation Letters*, pp. 3283–3290, 2024.
- [22] J. Nocedal and S. J. Wright, *Numerical optimization*. Springer, 2006.
- [23] P. F. Felzenszwalb and D. P. Huttenlocher, "Distance transforms of sampled functions," *Theory of Computing*, pp. 415–428, 2012.
- [24] B. Zhou, F. Gao, L. Wang, C. Liu, and S. Shen, "Robust and efficient quadrotor trajectory generation for fast autonomous flight," *IEEE Robotics and Automation Letters*, pp. 3529–3536, 2019.
- [25] Z. Wang, X. Zhou, C. Xu, and F. Gao, "Geometrically constrained trajectory optimization for multicopters," *IEEE Transactions on Robotics*, pp. 3259–3278, 2022.
- [26] D. C. Liu and J. Nocedal, "On the limited memory BFGS method for large scale optimization," *Mathematical Programming*, pp. 503–528, 1989.
- [27] A. S. Lewis and M. L. Overton, "Nonsmooth optimization via quasi-newton methods," *Mathematical Programming*, pp. 135–163, 2013.
- [28] C. Bai, T. Xiao, Y. Chen, H. Wang, F. Zhang, and X. Gao, "Faster-LIO: Lightweight tightly coupled LiDAR-inertial odometry using parallel sparse incremental voxels," *IEEE Robotics and Automation Letters*, pp. 4861–4868, 2022.
- [29] R. Khanam and M. Hussain, "YOLOv11: An overview of the key architectural enhancements," *arXiv preprint arXiv*, 2024.
- [30] L. Bertoni, S. Kreiss, and A. Alahi, "MonoLoco: Monocular 3D pedestrian localization and uncertainty estimation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 6861–6871.