

Flow-Aided Flight Through Dynamic Clutters From Point To Motion

Bowen Xu, Zexuan Yan, Minghao Lu, Xiyu Fan, Yi Luo, Youshen Lin,
Zhiqiang Chen, Yeke Chen, Qiyuan Qiao and Peng Lu

Abstract—Challenges in traversing dynamic clutters lie mainly in the efficient perception of the environmental dynamics and the generation of evasive behaviors considering obstacle movement. Previous solutions have made progress in explicitly modeling the dynamic obstacle motion for avoidance, but this key dependency of decision-making is time-consuming and unreliable in highly dynamic scenarios with occlusions. On the contrary, without introducing object detection, tracking, and prediction, we empower the reinforcement learning (RL) with single LiDAR sensing to realize an autonomous flight system directly from point to motion. For exteroception, a depth sensing *distance map* achieving fixed-shape, low-resolution, and detail-safe is encoded from raw point clouds, and an environment change sensing *point flow* is adopted as motion features extracted from multi-frame observations. These two are integrated into a lightweight and easy-to-learn representation of complex dynamic environments. For action generation, the behavior of avoiding dynamic threats in advance is implicitly driven by the proposed change-aware sensing representation, where the policy optimization is indicated by the relative motion modulated distance field. With the deployment-friendly sensing simulation and dynamics model-free acceleration control, the proposed system shows a superior success rate and adaptability to alternatives, and the policy derived from the simulator can drive a real-world quadrotor with safe maneuvers.

Index Terms—Aerial Systems; Perception and Autonomy; Reinforcement Learning; Autonomous Vehicle Navigation

I. INTRODUCTION

The past few decades have witnessed a rapid development of autonomous navigation, where elaborate mapping procedures and efficient trajectory representations [1] provide robust guidance for flight in complex environments and can be extended to swarm missions [2]. But their dependence on the perception level is still limited by compound errors introduced from inaccurate sensor observations and processing delays due to explicit environment representations. Learning-based methods attempt to alleviate these nuisances by utilizing data-driven approaches to draw complex decision-making knowledge into lightweight neural networks for high real-time performance flight through offline learning [3] [4]. Recent advances have revealed the potential for generating navigation behavior from

Manuscript received: June 7, 2025; Revised October 10, 2025; Accepted November 6, 2025. This paper was recommended for publication by Editor Ashis Banerjee upon evaluation of the Associate Editor and Reviewers' comments. This work was supported by General Research Fund under grant no. 17204222. (Bowen Xu and Zexuan Yan contributed equally to this work.) (Corresponding author: Peng Lu)

The authors are with the Adaptive Robotic Controls Lab (ArcLab), Department of Mechanical Engineering, The University of Hong Kong, Hong Kong SAR 999077, China. (email: lupeng@hku.hk).

Digital Object Identifier (DOI): see top of this page.

©2026 IEEE

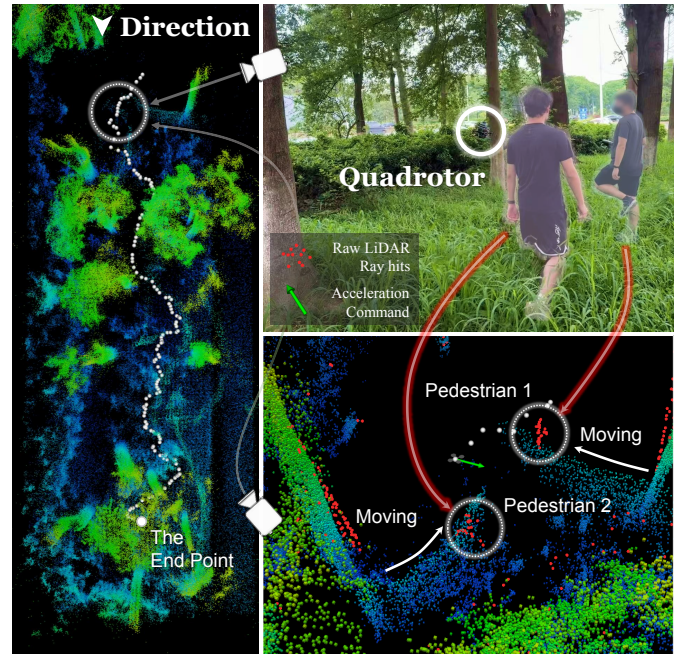


Fig. 1. Outdoor flight in natural clutters with pedestrians. The left part shows the executed trajectory of the neural controller. The snapshots on the right part demonstrate the avoidance response when confronted with both dynamic pedestrians and static trees. The movement of the pedestrians can be seen clearly in the “ghost trail” of the map built after the experiment.

perception to action, such as using 3D LiDAR point clouds [5], depth image [6], and monocular optical flow [7], which can be transferred to the real world at a low cost.

The presence of dynamic objects makes it difficult for the robot to determine the complete environment occupancy by compensating for ego-motion. In order to intuitively account for obstacle movement in the planning session, existing works choose to distinguish dynamic objects from the environment. However, occluded motion in onboard observations constitutes an important component of sensed environmental dynamics, but object-level understanding [8] requires consistent and coherent localized characteristics and motion models. This causes potential neglect of incomplete motions. Furthermore, the simultaneous state estimation of multiple objects [9] creates a huge computational burden, and sensor fusion strategies [10] exacerbate the perceptual latency. More importantly, the hand-crafted detection-tracking process requires multiple processes to work in concert, introducing inter-module compound errors [11]. The above factors make the dynamic perception of object-aware approaches untimely and unreliable as inputs

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

to subsequent modules, e.g., state predictions for trajectory optimization or network policies for action generation.

In this letter, we present an autonomous flight system that leverages LiDAR sensing in conjunction with a generalizable RL technique. Considering that using raw points for training is computationally inefficient, and oversampling leads to a loss of obstacle information, we design a low-resolution distance map with fixed shape and preserving detail safety based on the robot-centric nature of LiDAR scans. For dynamic environments, we aim to find a fine-grained representation that is correlated with environmental changes, and is consistent with the distance map in scale, shape, and positional meaning, which naturally guides us to supplement pixel motion for each distance sensing. We use this point-based relative motion feature **point flow** obtained from a pre-trained neural optical flow estimator, together with the depth perception **distance map** as the policy input. We implement the training by encouraging evasive maneuvers against dynamic obstacles that are implicitly driven by the change-aware input, which is correlated to the obstacle relative motion during policy optimization. The stable LiDAR-inertial proprioception, compact and realistically portable sensing simulation, and end-to-end policy architecture allow the system to be applied to physical aerial robots.

Our main contributions can be summarized as follows:

- *Propose an efficient LiDAR representation combining the depth and environment change sensing*, where a distance map achieves low-resolution depth perception while preserving detail safety, and a point flow describes the holistic relative motion of the surroundings.
- *Develop a generalizable RL training for dynamic obstacle avoidance* through learning early avoidance maneuvers from implicit environmental motion features, where the policy optimization is hinted by the relative motion-reshaped distance field of the dynamic obstacle.
- *Integrate a lightweight autonomous flight system from point to motion* without the need of object detection, tracking, and prediction. Extensive simulations and real-world experiments verify its higher success rate than baselines, especially in complex dynamic environments.

II. RELATED WORK

A. Hierarchical Object-aware Dynamic Obstacle Avoidance

Technical routes to discriminate dynamic objects from the environment for obstacle avoidance have been widely used. Some reactive solutions are based on repulsive potential field from moving object events [12] and trajectory from model predictive control (MPC) [13], but the in-the-moment optimization only focuses on the local optimum of safety and ignores the long-term requirements of navigation tasks such as smooth movement and heading to the goal, resulting in low adaptability and efficiency in complex dynamic environments. More complete systems use both dynamic clusters and static occupancy for trajectory optimization, where rgb-d based system [8] projects small object recognition results into the depth domain and point cloud based systems [14] [9] cluster and segment the spatial information, these methods estimate

the states of dynamic obstacles and further predict their state sequences for calculating the cost function of planning. However, the comprehensive modeling of dynamic objects introduces compound errors for the final prediction module, which leaves the trajectory optimization module to rely on inaccurate environmental references. Building on these efforts, learning from object recognition is adopted to bypass the unreliable prediction-planning pipeline, where [10] uses the encoded detection result image as pedestrian information and [11] utilizes obstacle state vectors as dynamism sensing. But fundamentally limited by the object-aware strategy, their detection-tracking modules underutilize the environment dynamics for struggling to account for the incomplete motions caused by occlusions, and their real-time performance decreases drastically when the quantity of dynamic obstacles increases, thus failing to maintain safety in dynamic and cluttered scenarios. In comparison, our method does not scrutinize the dynamic-static distinction in sensing, where the environmental changes throughout the observation can be uniformly represented at an object-free pixel level for end-to-end evasive maneuvers optimization.

B. End-to-end Navigation in Dynamic Clutters

For dynamic environment navigation that maps perception to action, existing approaches leverage the high-frequency inference properties of neural controllers for adaptability to changing environments, where [15] learns with differentiable simulations using low-resolution depth images, and [16] derives velocity and steering angle from single-frame LiDAR scans. These works rely on static sensing that achieves avoidance of simple moving objects, but is difficult to generalize to complex dynamic environments. This is because the lack of historical information prevents them from sensing additional threats introduced by obstacle motion. To enable dynamic awareness, learning from past sensing is presented. Velocity command of [17] is mapped from the last three frames of LiDAR scans, and acceleration control of [18] is achieved with an obstacle image stitched from past scan vectors. However, motion patterns in multi-frame observations are difficult to capture without explicit supervision. This lack of emphasis on inter-observation temporal variations makes it difficult to learn safety maneuvers triggered by obstacle movement patterns, which can lead to overfitting of the training dynamic scenes. On the basis of the above works, integrating multi-frame perception is proposed to obtain more explicit motion-related inputs, where [19] represents obstacle kinematics from point cloud alignment results and [20] denotes scene temporal evolutions as long short-term memory-encoded depth images. For learning avoidance behaviors, these representations are of good benefit, but obtaining them requires specially optimized feature extraction modules. Training these modules from scratch entails expensive customized datasets, whereas co-optimization in RL leads to unstable and unnecessary updating of the dynamic perception module subject to policy gradients, making it difficult to learn generic features. In contrast, our approach utilizes pre-trained weights to align the geometric structure of historical observations with the semantics of pixel motion, obtaining a low-noise representation which is highly correlated

with the subtask space of environmental changes. The flow attention serves as an aid to depth perception, providing a reliable environment representation which is further embedded to guide navigation decisions.

III. FRAMEWORK

A. Distance Map Encoding

Given the LiDAR horizontal and vertical field of view (FOV) and their boundaries as $F_H \in (F_{Hh}, F_{Hl})$, $F_V \in (F_{Vh}, F_{Vl})$, we consider all point clouds within the truncation distance d_{\max} . We define a current scanning point set in the body frame as ${}_B P$, and for each point $p_i \in {}_B P$, it has a coordinate (x_i, y_i, z_i) and a distance to the quadrotor $d_i = \sqrt{x_i^2 + y_i^2 + z_i^2}$. A raycast operation is first applied to organize the raw point cloud into an ordered point set with a fixed shape of $r_h \times r_v \times 3 = 108 \times 18 \times 3$, where space within the FOV is divided by intervals of $i_h = F_H/r_h, i_v = F_V/r_v$ in the horizontal and vertical direction. Each point can determine which partition it belongs to according to its azimuth θ_h and elevation θ_v relative to the quadrotor. The point within the s th (m th in horizontal, n th in vertical, $s = r_h(n-1) + m$) partition is $P_s = \{p_i \in {}_B P \mid \xi\}$.

$$\theta_{hi} = \tan^{-1}(y_i/x_i), \theta_{vi} = \sin^{-1}(z_i/d_i),$$

$$\xi := \frac{\theta_{hi} - F_{Hl}}{i_h} \in (m-1, m), \frac{\theta_{vi} - F_{Vl}}{i_v} \in (n-1, n). \quad (1)$$

For the points in each partition, we take the point with the smallest distance value as the representative point of this partition. This approach retains more obstacle details than simply downsampling, essentially inflating the nearest obstacle point in its partition to improve safety. The point p_s corresponding to the s th partition in the raycast point set can be selected as $p_s = \arg \min_{p_i \in P_s} d_i$.

Then, to obtain the image-form depth representation, we first remove unnecessary detections from the ceiling or ground, only keeping the point cloud within a relative height interval of the quadrotor. Then, we take the distance value of each partition of the raycast point set to encode a high-resolution (108×18) grayscale image. A further minimal pooling is adopted to get the low-resolution (36×6) distance map.

B. Point Flow Estimation

To enable parallel flow computation in training, and inference acceleration in testing, we use a learning-based method, NeuFlowV2 [21], to provide real-time and robust flow estimation for our system. As shown in Fig. 2, we resize the grayscale image encoded from the raycast point set from (108×18) to (96×16) using bilinear interpolation, and combine three consecutive frames into a single 3-channel image to fit the input of the pre-trained Neuflow checkpoint. Then, a sliding window with length n is applied to ensure continuous and high-frequency input. We consider the previous concatenated image I_{k-n+3} , which is $n-3$ frames ahead of the current concatenated image I_k . The previous frame pixels (u_{k-n+3}, v_{k-n+3}) are re-projected to the current frame pixels (u_k, v_k) after relative motion with the quadrotor. The real flow F_k can be estimated as $\hat{F}_k = \text{neuflow}(I_{k-n+3}, I_k)$ and

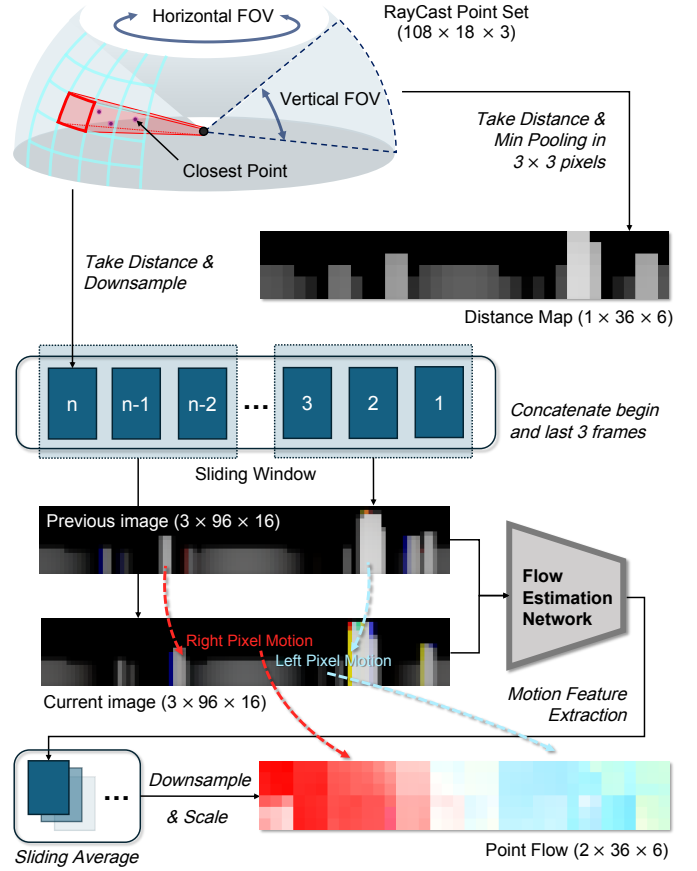


Fig. 2. Illustration of the change-aware LiDAR sensing representation. The distance map is encoded by robot-centric omnidirectional depth perception, which ensures safe sensing of obstacle details through nearest selection. The point flow is derived from pseudo-RGB image pairs stacked with 3 consecutive grayscale images, which are fed into a flow estimation module that does not update during RL training for motion feature extraction.

a sliding average is carried out on 5 frames of continuous flow with a window size of $5/50 \text{ Hz} = 0.1 \text{ s}$ to balance the estimation error.

After downsampling the average flow $\bar{F}_{k \rightarrow k-4}$ to the low-resolution form $\bar{F}'_{k \rightarrow k-4}$ to match the size of the distance map, we can obtain the point flow by further scaling:

$$\bar{F}'_{k \rightarrow k-4} = (\bar{\Delta}u'_{k \rightarrow k-4}, \bar{\Delta}v'_{k \rightarrow k-4})_{2 \times 36 \times 6},$$

$$\text{pointflow}_k = \left(\frac{\bar{\Delta}u'_{k \rightarrow k-4}}{z_H}, \frac{\bar{\Delta}v'_{k \rightarrow k-4}}{z_W} \right)_{2 \times 36 \times 6}, \quad (2)$$

where the horizontal and vertical flow values of $\bar{F}'_{k \rightarrow k-4}$ are scaled by z_H and z_W respectively to re-balance the attention on the pixel motion of different directions.

C. System Overview of Observation and Action

The framework of the system is shown in Fig. 3. The distance map and the point flow are concatenated in the channel dimension into a tensor of shape $(3, 36, 6)$ for working together as the LiDAR sensing representation. The unit vector pointing to the goal \mathbf{g} , the velocity \mathbf{v} , and the last action \mathbf{a}_{last} are all 3D vectors with the shape of $(3,)$, which inform the target direction and ego-motion. A network policy is used for

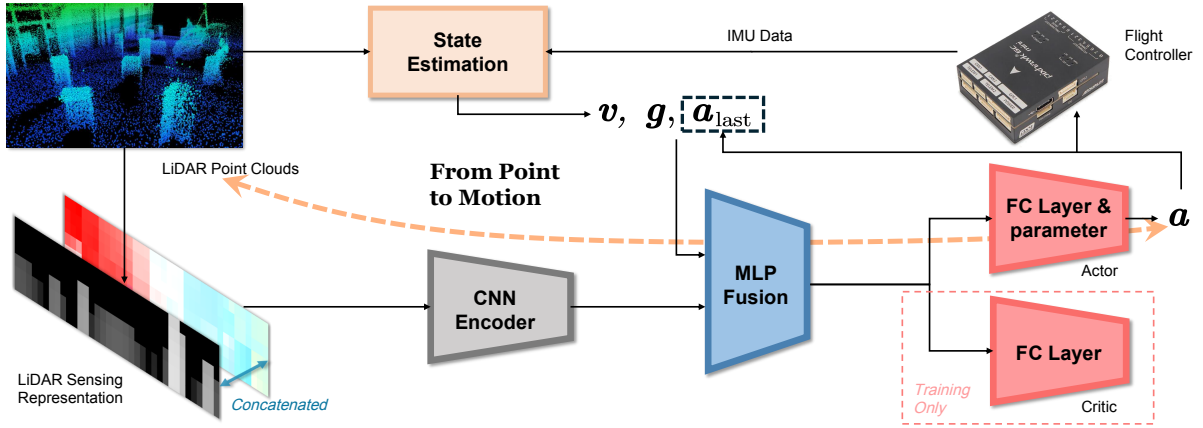


Fig. 3. Implementation of the autonomous flight system from point to motion. State estimation [22] of the quadrotor is realized through the collaborative use of the LiDAR point clouds and the IMU data, which is a robust real-world substitute for ground truth states used in training. The unit vector pointing to the goal \mathbf{g} , the velocity \mathbf{v} and the last output \mathbf{a}_{last} form the state inputs of the system, which are used to feed into the MLP fusion module together with the LiDAR sensing representation features extracted by the CNN encoder. The output of the system is the desired acceleration command \mathbf{a} .

acceleration command inference of (3). These system designs select both observation and action to be at the kinematic level, thereby achieving a universal navigation method enabling policy deployment on different quadrotors across training and reality, regardless of the inconsistent underlying dynamics parameters.

IV. REINFORCEMENT LEARNING FOR THE OBSTACLE AVOIDANCE POLICY

A. Policy Representation

We employ a partially shared architecture to model the policy and the value functions, which share an encoder based on convolutional neural network (CNN) and fully connected (FC) layer. The convolutional part has three layers, whose structure [output channel, kernel size, stride, padding] is respectively [4, (5,3), (2,1), 0], [16, (5,3), (2,1), (2,1)] \times 2. The output of the CNN layer is flattened and is then fed into the FC layer of 128 units and normalized for extracting LiDAR features. After being concatenated with state inputs, the observations are fused by the multi-layer perceptron (MLP) of [256, 256] and fed into an actor-critic module.

B. Reward Function

The reward function is designed according to the safety-assured, goal-oriented, and dynamically feasible requirements of navigation tasks, and a dynamic obstacle avoidance reward is adopted for better adaptation to the dynamic environment. The reward function r can be described as:

$$r = r_{\text{state}} + r_{\text{goal}} + r_{\text{safety}} + r_{\text{dobs}}, \quad (3)$$

where r_{state} , r_{goal} , r_{safety} and r_{dobs} represent the state reward, goal reward, safety reward and dynamic obstacle avoidance reward.

1) *State reward*: The quadrotor states are softly constrained to the desired range by a limiting function:

$$f_l(s, \beta_s, s_{\text{low}}, s_{\text{high}}) = \log(e^{-\beta_s(\max(s_{\text{low}} - s, 0) + \max(s - s_{\text{high}}, 0))} + 1), \quad (4)$$

where s , β_s , and $[s_{\text{low}}, s_{\text{high}}]$ represent the quadrotor state magnitude, the corresponding discount factor, and the desired range. Then, we define a velocity reward r_v , an acceleration reward r_a , a height reward r_h to encourage the desired flight states, and a jerk reward r_j to promote a smooth trajectory.

$$\begin{aligned} r_v &= f_l(\|\mathbf{v}\|, \beta_v, v_{\min}, v_{\text{ref}}), \quad r_a = f_l(\|\mathbf{a}\|, \beta_a, a_{\min}, a_{\text{ref}}), \\ r_h &= f_l(p_z, \beta_h, h_{\min}, h_{\max}), \quad r_j = 1/(1 + \|\mathbf{a}_t - \mathbf{a}_{t-1}\|), \quad (5) \\ r_{\text{state}} &= \lambda_v r_v + \lambda_a r_a + \lambda_h r_h + \lambda_j r_j, \end{aligned}$$

where \mathbf{v} and \mathbf{a} are the quadrotor velocity and the acceleration command. v_{ref} (a_{ref}) and v_{\min} (a_{\min}) are the reference and minimum velocity (acceleration), which define the expected and minimum-acceptable agility. p_z is the z coordination of the quadrotor position \mathbf{p} , h_{\min} and h_{\max} are the bounds of the appropriate flight height, and λ is the relative weight.

2) *Goal reward*: The behavior of flying to the goal is guided from two perspectives: the projection of the velocity in the goal direction r_{dir} and the change of distance to the goal r_{dis} , where r_{dir} encourages a goal-aligned velocity direction and r_{dis} encourages an actual flight progress.

$$\begin{aligned} r_{\text{dir}} &= \min(\mathbf{v} \cdot \mathbf{p}_g / \|\mathbf{p}_g\|, v_{tr}), \quad r_{\text{dis}} = e^{d_g(t-1) - d_g(t)} - 1, \quad (6) \\ r_{\text{goal}} &= \lambda_{\text{dir}} r_{\text{dir}} + \lambda_{\text{dis}} r_{\text{dis}}, \end{aligned}$$

where \mathbf{p}_g and d_g are the goal position and the distance to the goal. The truncated projection velocity $v_{tr} = 0.4v_{\text{ref}}$ promotes a later training stage that shifts the dominance of r_{goal} from the saturated r_{dir} to r_{dis} , towards a safety-balanced goal-directed skill than the initial velocity alignment.

3) *Safety reward*: Safety reward is designed as a two-stage reward that can be derived directly from LiDAR scans.

$$\begin{aligned} d_s &= \begin{cases} \text{avg}\{d_{li} - r_d \mid d_{li} < d_{th}\}, & \min(d_{li}) < d_{th} \\ \min(d_{li} - r_d), & \text{otherwise} \end{cases} \quad (7) \\ r_{\text{safety}} &= \lambda_s \max(\log(d_s), \log(d_{tr})), \end{aligned}$$

where d_{li} and r_d are the distance of the i th ray hit and the radius of the quadrotor, $d_{th} = 1$ is the distance threshold. The truncated logarithmic distance $\log(d_{tr}) = -5$ means that

when the agent is close to the obstacle, extreme negative values should be avoided to dominate the total reward.

4) *Dynamic obstacle avoidance reward*: The dynamic obstacle avoidance reward is considered on a horizontal level and realized by reshaping the distance field of the dynamic obstacles, which aims to encourage the quadrotor to avoid potential danger in advance. The extra collision threat region of a dynamic obstacle is defined by the following condition.

$$\theta = \cos^{-1}\left(\frac{(\mathbf{p} - \mathbf{p}_{\text{dobs}}) \cdot (\mathbf{v}_{\text{dobs}} - \mathbf{v})}{\|\mathbf{p} - \mathbf{p}_{\text{dobs}}\| \cdot \|\mathbf{v}_{\text{dobs}} - \mathbf{v}\|}\right), \quad (8)$$

where \mathbf{p}_{dobs} and \mathbf{v}_{dobs} are the position and velocity of a dynamic obstacle. Eq. (8) evaluates the angle between the relative velocity (obstacle velocity in the quadrotor body frame) and the connection between the obstacle and the quadrotor. Then the reshaping coefficient k can be defined with θ as follows:

$$k_v = \|\mathbf{v}_{\text{dobs}} - \mathbf{v}\|, \quad k_\theta = 1 - \left(\frac{2\theta}{\pi}\right), \quad k_d = e^{\frac{1}{1+d_v}}, \quad (9)$$

$$k = 1 + k_v \cdot k_\theta \cdot k_d \quad \text{if } \theta \in (0, \frac{\pi}{2}), \quad \text{otherwise } k = 1,$$

where d_v represents the distance from the quadrotor to the obstacle velocity line, and k is used to discount the original distance field as in Fig. 4.

Considering only the dynamic obstacles hit by the rays is reasonable, but this will make the reward highly random due to the fast-changing nature of dynamic surroundings. One way to stabilize it is to relax the hitting constraint and replace it with considering all dynamic obstacles within a certain distance to the quadrotor. Although privileged information is introduced, this solution can help better convergence and can be easily generalized to scenarios with occlusions.

$$r_{\text{dobs}} = \begin{cases} \frac{\lambda_d}{m} \sum_m \max(\log(\frac{d_{\text{dobs}_i} - r_d}{k_i}), \log(d_{tr})), & m > 0 \\ \max(r_{\text{dobs}}(t-1), \dots, r_{\text{dobs}}(0)), & \text{otherwise} \end{cases} \quad (10)$$

where d_{dobs_i} is the distance between the i th dynamic obstacle and the quadrotor, m is the number of dynamic obstacles within $0.75d_{\text{max}}$. When no dynamic obstacles are in this range, the reward is set to the maximum of the previously occurring r_{dobs} to encourage this behavior of totally away from dynamic threats, while reducing the reward sparsity.

C. Training Setup and Implementation

As shown in Fig. 5, static columns and walls are set up to guide 3D avoidance actions. Dynamic obstacles are set to move in uniform linear motions with a randomly assigned velocity within $[1, 5] \text{ m/s}$. We define early termination for (i) agent exceeding the limit of velocity ($1.2v_{\text{ref}}$), acceleration ($1.5a_{\text{ref}}$) or height, (ii) being too close to obstacles, and (iii) straying too far from the flight course. The max episode length is set as higher than the process buffer required for flying to the goal, and r_{dis} and r_v are cut off when the quadrotor is close enough to the goal, changing the training mode from a navigation task to a safe wandering around the goal. The implementation of the underlying proximal policy optimization (PPO) [23] training toolchain is highly inspired by Omnidrones [24], which is built on Nvidia Isaac Sim.




Fig. 4. Effects of reshaping the distance field of dynamic obstacles. The distance field is shown in a range of $5 \times 5 \text{ m}^2$ when the dynamic obstacle is at the same orthogonal velocity of 1 m/s as the quadrotor. This exhibits the extra threat given to $\theta \in (0, \pi/2)$ is roughly symmetrical about the relative velocity $\mathbf{v}_{\text{dobs}} - \mathbf{v}$ and skewed toward the absolute obstacle velocity \mathbf{v}_{dobs} .

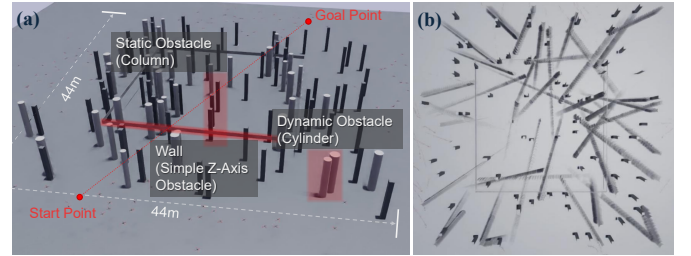


Fig. 5. Training Environment and Obstacle Setup. (a) Parallel training in 128 environments with centrosymmetrically set start-goal points, and walls are set at the same height as the start-goal points. (b) Top-down view of the trajectories of the dynamic obstacles performing uniform linear motions.

V. EXPERIMENTS AND EVALUATIONS

A. Simulation Experiments

1) *Benchmark with previous systems*: We benchmark our system against recent representative works: a hierarchical rule-based method FAPP [9], a hierarchical learning-based method NavRL [11], and an end-to-end method Obsnet [18], all of which serve as the state-of-the-art navigation approaches. For a reasonable comparison, we set the reference acceleration to 10 m/s^2 for acceleration control-based [18] and the proposed method, and the acceleration limit of [9] to 10 m/s^2 . In terms of velocity, [18] does not need a reference velocity input; the proposed method accepts v_{ref} of 5 m/s ; and the local greedy trajectory optimization design of [9] requires a more stringent velocity limit, which is set to 3.5 m/s . For [11] trained towards low-speed scenarios, its velocity command output is set within $[0.5, 1.5] \text{ m/s}$, and its perception range is also reduced to 5 m compared to 10 m for other methods to match its agility.

The simulation environment is shown in Fig. 6, where the dynamic obstacles perform variable-speed curve motion, which is unseen compared to that in training. We counted the success rate η from 20 experiments and other performance metrics from 5 successful flights. The evaluation of performance metrics is performed on an i5-14600KF CPU and RTX4070Ti GPU, including the average flight speed v_a , planning latency from perception to action t_p , path efficiency $R_l = l/\|\mathbf{p}_g - \mathbf{p}_s\|$ meaning the path length l divided by the distance from start \mathbf{p}_s to goal \mathbf{p}_g , and safety distance d_s .

The perception modules in [9] and [11] work well in low obstacle densities. Still, the increased computation burden

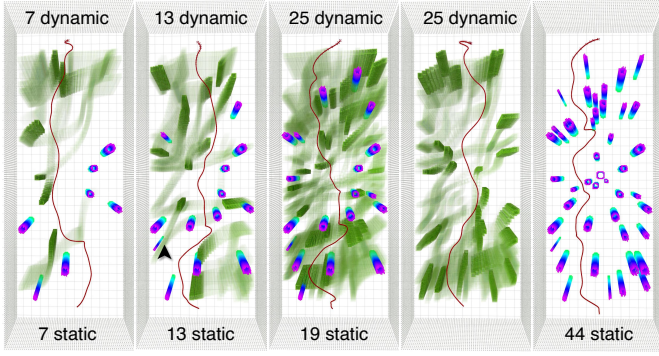


Fig. 6. Simulated flight slices of the proposed system. The benchmark simulation experiments are executed in a $25 \times 10 \text{ m}^2$ field, where static obstacles contain both thick and thin columns, and the maximum speed of dynamic obstacles is limited to below 5 m/s .

TABLE I
BASELINE COMPARISON IN CLUTTERED ENVIRONMENTS WITH DIFFERENT DENSITIES OF STATIC AND DYNAMIC OBSTACLES

Scenario	Method	$\eta(\%)$	$v_a(\text{m/s})$	$t_p(\text{ms})$	R_i	$d_s(\text{m})$
7 dynamic and 7 static	Ours	95	3.14	9.58	1.13	2.80
	[9]	90	2.20	9.31	1.08	2.43
	[11]	45	0.88	24.30	1.04	1.71
	[18]	20	5.73	4.05	1.05	0.97
13 dynamic and 13 static	Ours	65	2.55	10.38	1.43	2.05
	[9]	30	1.88	12.75	1.22	1.93
	[11]	30	0.89	26.19	1.06	1.62
	[18]	30	5.61	5.06	1.15	0.90
25 dynamic and 19 static	Ours	40	2.29	12.29	1.20	1.83
	[9]	25	2.09	21.36	1.21	1.74
	[11]	20	0.90	33.85	1.20	1.39
	[18]	20	5.35	6.92	1.08	0.93
25 dynamic	Ours	50	2.92	10.26	1.27	2.30
	[9]	20	1.94	17.01	1.31	2.05
	[11]	15	0.87	28.28	1.25	1.48
	[18]	25	5.33	4.32	1.07	0.96
44 static	Ours	60	2.15	8.88	1.24	1.52
	[9]	75	1.84	22.89	1.07	1.36
	[11]	80	0.92	27.09	1.13	1.32
	[18]	/	/	/	/	/

causes significant latency when the scene complexity rises. In contrast, the proposed method and [18], due to the fixed-shape sensing representation, their real-time performance is only slightly affected by the number of raw point clouds. The online optimization of [9] enhances the safety by explicitly pushing the trajectories to the free space. But a large number of motion-variable obstacles elicit lagging state estimates, while frequent replanning is triggered but relies only on predictions discredited by compound errors. This results in unsafe behaviors in highly dynamic settings. The learning-based object-aware approach [11] faces a similar dilemma. The sacrifice of agility realizes higher path efficiency and excellent static navigation performance. However, attributed to the strategy of learning avoidance maneuvers from dynamic-static compositions of the environment, its hyper-efficient and adaptive RL network policy is held back by its heavy detection-tracking modules and extra safety actions. This fundamentally leads to disadvantages in dynamic scenarios.

Meanwhile, insights that discard intermediate modules show




Fig. 7. Reward components that show outstanding changes after ablation. The curves show rewards of the ablated systems after training the same steps as the proposed system, including the (a) velocity and (b) jerk rewards. They demonstrate the extent to which the agent is in the desired velocity interval and the extent to which the acceleration command mutates.

TABLE II
ABLATION STUDY ON POINT FLOW AND DYNAMIC OBSTACLE AVOIDANCE REWARD WITH DIFFERENT OBSTACLE SPEEDS

Obstacle Speed	Ablation	$\eta(\%)$	α_c	$v_a(\text{m/s})$
Low ($\ \mathbf{v}_{\text{dobs}}\ < 1 \text{ m/s}$)	Ours	90	/	3.09
	w/o flow	60	0.67	2.87
	w/o r_{dobs}	70	0.78	3.73
Medium ($\ \mathbf{v}_{\text{dobs}}\ < 2 \text{ m/s}$)	Ours	80	/	3.08
	w/o flow	50	0.63	2.68
	w/o r_{dobs}	55	0.69	3.57
High ($\ \mathbf{v}_{\text{dobs}}\ < 3 \text{ m/s}$)	Ours	60	/	2.92
	w/o flow	35	0.58	2.51
	w/o r_{dobs}	40	0.67	3.38

a high adaptability to moving objects. [18] learns directly from stacked historical scanning sequences, which can traverse dynamic clutters at high speeds. However, the lack of awareness of motion patterns between frames leads to difficulties in capturing threat differences across obstacle situations. This further gives rise to low safety and overfitting to dynamic scenarios, as well as breaking down in static scenarios. As can be seen from Table I, thanks to the change-aware LiDAR sensing representation and the flight skills empowered by relative motion-guided policy optimization, our method exhibits the most reliable safety in the vast majority of scenes, maintaining a success rate of 40% in the most complex environment containing both dynamic and static obstacles compared to the less than 30% of the other methods. Furthermore, in comparison with [18], although the agility of our method is not as high, it has the advantage of being adaptively regulated by the environment complexity based on the reference input.

2) *Ablation on method components*: We perform ablations on the point flow and the dynamic obstacle avoidance reward r_{dobs} , and the ablated training results are shown in Fig. 7. We conduct experiments in dense scenarios containing 33 dynamic obstacles and implement control groups with different obstacle speed limits, which helps to directly assess the ablation impact on dynamic obstacle avoidance performance. The results are shown in Table II, where the newly introduced metric α_c represents the ratio of the success rate of the ablated system to the success rate of the proposed system.

Leveraging the intuitive depth sensing, the system with ablated point flow still ensures the basic safety, but the loss of relative motion perception means it can only keep reacting to the dynamic obstacles under the static assumptions. The

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.



Fig. 8. Real-world experiments. (a) Traversal through a box environment. (b) Agile flight with high reference acceleration input. (c) Frequently interacting with pedestrians. (d) Highlight of a backward avoidance when performing a similar experiment to (c) in a lower agility. (e) Reacting to fast-moving objects. (f) Navigation among static and dynamic obstacles. (g) Details of adapting to an obstacle emergency stop in experiment (f).

difficulty in taking efficient and early avoidance actions directly leads to the decrease of flight speed and the reduction of velocity reward during training. The system without dynamic obstacle avoidance reward can not effectively correlate the estimated environmental changes with the desired behavior. It struggles to realize the differences in threat prioritization when handling a multi-dynamic obstacle scenario, thus it prefers to give over-excited maneuvers with sudden changes to achieve avoidance when the safety distance is very small. This causes a drop in jerk reward and also results in an inappropriately excessive agility. In a general view, the effects of the ablated system components on flight safety become more obvious as the dynamics of the environment increase, which are manifested by the gradual decrease of α_c , which reflects the proximity in reliability.

B. Real-World RL Deployment

Thanks to the kinematic-level observation and action designs, we can use a physical quadrotor that has different dynamics parameters from training. We apply the policy obtained from the sensor attitude-ignored training to the rigidly connected LiDAR-quadrotor in reality, where the Livox Mid-360 LiDAR is used to provide point cloud perception and the PX4 flight controller is used to track the desired acceleration as in Fig. 3. We use TensorRT for inference acceleration on the on-board computing unit Nvidia Jetson Orin NX, which ensures a control frequency of 50 Hz.

C. Real-World Experiments

1) *Traversal through a box environment & Agile flight with high reference acceleration input:* To evaluate the obstacle avoidance ability in static environments, we design a

scene composed of boxes placed horizontally or vertically at different heights, which is similar to the column obstacle in training. As shown in Fig. 8(a), the quadrotor can make adaptive decisions to fly through the free space on either side or above the obstacle. In order to test the ability of agile flight, we retain a few columnar obstacles and increase the expected agility to a high level. As shown in Fig. 8(b), the quadrotor can nimbly shuttle among obstacles and reach a maximum speed of 14.57 km/h (4.05 m/s).

2) *Frequently interacting with pedestrians & Reacting to fast-moving objects:* To assess the ability of handling dynamic objects, we arrange for multiple pedestrians to walk in chaotic and subjectively varying trajectories and speeds, which is different from the uniform linear motion of dynamic obstacles in training. As shown in Fig. 8(c), the quadrotor can efficiently fly through areas without potential dangers. When the threat of the dynamic environment is significant, the policy does not greedily insist on the fastest-forward maneuvers. Instead, it chooses actions such as deceleration or stopping to ensure safety. In a similar environment but lower expected agility setup, backward action can also be triggered to achieve a low-speed avoidance solution, as shown in Fig. 8(d). In order to test the response to fast-moving obstacles, we throw boxes along the path ahead of the quadrotor to interfere with its normal flight, as shown in Fig. 8(e), which is similar to the scenario of a fast-moving obstacle appearing from the side in training. The quadrotor can continuously respond to multiple fast-moving thrown obstacles and always maintain a safe distance from them.

3) *Navigation among static and dynamic obstacles & Flying in natural clutters with pedestrians:* To evaluate the navigation ability in a complex dynamic environment, we design a scene containing columns and pedestrians. The dense occlusion

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

and unpredictable dynamic behaviors pose challenges for perception and planning. As shown in Fig. 8(f), the quadrotor repeatedly adjusts its maneuverability and travel direction to safely bypass the dynamic and static obstacles. Meanwhile, when the speed of the dynamic obstacle undergoes a sudden change, such as an emergency stop in Fig. 8(g), the policy can still handle it robustly without being triggered to take excessive actions. We also deploy the system among unevenly distributed trees while integrating pedestrians into the environment. The challenges posed by natural clutters primarily stem from the domain gap between training, including highly unstructured observations, a significant amount of uninformative points, and unaccounted-for wind disturbances. As shown in Fig. 1, in this long-distance task, the policy can be easily generalized to tree scenes and navigate flexibly through pedestrians and trees.

D. Discussions on Robustness and Limitations

The robustness of the system is reflected in the performance under sensor degradation and unseen motion patterns. (i) The point flow yields reliable estimates under noise and sensor occlusions, whereas the degraded distance map may affect policy performance. (ii) Owing to the immediate and relative perception nature of the point flow, the policy can achieve robust near-distribution generalization to unseen obstacle motions with novel trajectories and speeds.

The main limitation of the system is reflected in some failure cases. (i) Due to the limitation of the sensor FOV and training obstacle setting, avoiding obstacles from above is not always successful. (ii) Due to the loss of dynamic perception caused by occlusion, collisions may happen with the high-speed obstacle that emerges beside others. (iii) Due to insufficient training sample coverage, collisions may happen with the obstacles with abrupt relative motion trend changes. In future work, we will incorporate risk awareness to overcome out-of-view and occluded threats, and supplement urgent maneuvers learning by diversifying training conditions.

VI. CONCLUSION

In this work, we showed that combining single LiDAR sensing with RL can realize an end-to-end autonomous flight system with effortless sim-to-real transfer. We demonstrated that, without introducing object detection, tracking, and prediction modules, lightweight access to complex dynamic environments can be achieved through the encoding of point cloud data and the extraction of multi-frame motion features. And we exhibited that, without trajectory generation, the dynamic obstacle avoidance behaviors can be implicitly driven by the relative motion-guided policy optimization.

REFERENCES

- [1] Z. Wang, X. Zhou, C. Xu, and F. Gao, "Geometrically constrained trajectory optimization for multicopters," *IEEE Transactions on Robotics*, vol. 38, no. 5, pp. 3259–3278, 2022.
- [2] X. Zhou, X. Wen, Z. Wang, Y. Gao, H. Li, Q. Wang, T. Yang, H. Lu, Y. Cao, C. Xu *et al.*, "Swarm of micro flying robots in the wild," *Science Robotics*, vol. 7, no. 66, p. eabm5954, 2022.
- [3] A. Loquercio, E. Kaufmann, R. Ranftl, M. Müller, V. Koltun, and D. Scaramuzza, "Learning high-speed flight in the wild," *Science Robotics*, vol. 6, no. 59, p. eabg5810, 2021.
- [4] Y. Song, K. Shi, R. Penicka, and D. Scaramuzza, "Learning perception-aware agile flight in cluttered environments," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 1989–1995.
- [5] G. Xu, T. Wu, Z. Wang, Q. Wang, and F. Gao, "Flying on point clouds with reinforcement learning," *arXiv preprint arXiv:2503.00496*, 2025.
- [6] J. Lu, X. Zhang, H. Shen, L. Xu, and B. Tian, "You only plan once: A learning-based one-stage planner with guidance learning," *IEEE Robotics and Automation Letters*, 2024.
- [7] Y. Hu, Y. Zhang, Y. Song, Y. Deng, F. Yu, L. Zhang, W. Lin, D. Zou, and W. Yu, "Seeing through pixel motion: Learning obstacle avoidance from optical flow with one camera," *IEEE Robotics and Automation Letters*, 2025.
- [8] M. Lu, H. Chen, and P. Lu, "Perception and avoidance of multiple small fast moving objects for quadrotors with only low-cost rgbd camera," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11 657–11 664, 2022.
- [9] M. Lu, X. Fan, H. Chen, and P. Lu, "Fapp: Fast and adaptive perception and planning for uavs in dynamic cluttered environments," *IEEE Transactions on Robotics*, 2024.
- [10] Z. Xie and P. Dames, "Drl-vo: Learning to navigate through crowded dynamic scenes using velocity obstacles," *IEEE Transactions on Robotics*, vol. 39, no. 4, pp. 2700–2719, 2023.
- [11] Z. Xu, X. Han, H. Shen, H. Jin, and K. Shimada, "Navrl: Learning safe flight in dynamic environments," *IEEE Robotics and Automation Letters*, 2025.
- [12] D. Falanga, K. Kleber, and D. Scaramuzza, "Dynamic obstacle avoidance for quadrotors with event cameras," *Science Robotics*, vol. 5, no. 40, p. eaaz9712, 2020.
- [13] Z. Xu, D. Deng, Y. Dong, and K. Shimada, "Dpmpc-planner: A real-time uav trajectory planning framework for complex static environments with dynamic obstacles," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 250–256.
- [14] Y. Wang, J. Ji, Q. Wang, C. Xu, and F. Gao, "Autonomous flights in dynamic environments with onboard vision," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 1966–1973.
- [15] Y. Zhang, Y. Hu, Y. Song, D. Zou, and W. Lin, "Learning vision-based agile flight via differentiable physics," *Nature Machine Intelligence*, pp. 1–13, 2025.
- [16] C. Pérez-D'Arpino, C. Liu, P. Goebel, R. Martín-Martín, and S. Savarese, "Robot navigation in constrained pedestrian environments using reinforcement learning," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 1140–1146.
- [17] T. Fan, P. Long, W. Liu, and J. Pan, "Distributed multi-robot collision avoidance via deep reinforcement learning for navigation in complex scenarios," *The International Journal of Robotics Research*, vol. 39, no. 7, pp. 856–892, 2020.
- [18] X. Fan, M. Lu, B. Xu, and P. Lu, "Flying in highly dynamic environments with end-to-end learning approach," *IEEE Robotics and Automation Letters*, 2025.
- [19] J. de Heuvel, X. Zeng, W. Shi, T. Sethuraman, and M. Bennewitz, "Spatiotemporal attention enhances lidar-based robot navigation in dynamic environments," *IEEE Robotics and Automation Letters*, 2024.
- [20] D. Hoeller, L. Wellhausen, F. Farshidian, and M. Hutter, "Learning a state representation and navigation in cluttered and dynamic environments," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5081–5088, 2021.
- [21] Z. Zhang, A. Gupta, H. Jiang, and H. Singh, "Neuflow v2: High-efficiency optical flow estimation on edge devices," *arXiv preprint arXiv:2408.10161*, 2024.
- [22] 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.
- [23] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [24] B. Xu, F. Gao, C. Yu, R. Zhang, Y. Wu, and Y. Wang, "OmniDrones: An efficient and flexible platform for reinforcement learning in drone control," *IEEE Robotics and Automation Letters*, vol. 9, no. 3, pp. 2838–2844, 2024.