

# SPOT: Spatio-Temporal Obstacle-free Trajectory Planning for UAVs in unknown dynamic environments

Astik Srivastava, Thomas J Chackenkulam, Bitla Bhanu Teja, Antony Thomas, and Madhava Krishna

**Abstract**—We address the problem of reactive motion planning for quadrotors operating in unknown environments with dynamic obstacles. Our approach leverages a 4-dimensional spatio-temporal planner, integrated with vision-based Safe Flight Corridor (SFC) generation and trajectory optimization. Unlike prior methods that rely on map fusion, our framework is mapless, enabling collision avoidance directly from perception while reducing computational overhead. Dynamic obstacles are detected and tracked using a vision-based object segmentation and tracking pipeline, allowing robust classification of static versus dynamic elements in the scene. To further enhance robustness, we introduce a backup planning module that reactively avoids dynamic obstacles when no direct path to the goal is available, mitigating the risk of collisions during deadlock situations. We validate our method extensively in both simulation and real-world hardware experiments, and benchmark it against state-of-the-art approaches, showing significant advantages for reactive UAV navigation in dynamic, unknown environments.

## I. INTRODUCTION

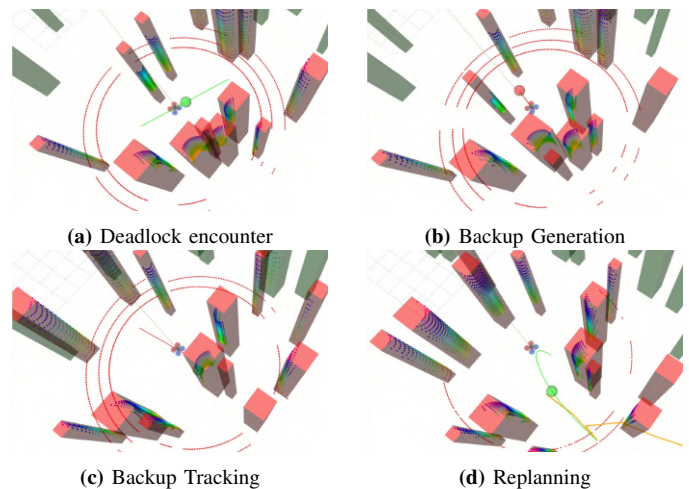
Quadrotors have emerged as the preferred unmanned aerial vehicle (UAV) platform for both civilian and military applications due to their simple mechanical design and ease of control. This has led to their widespread adoption in diverse domains such as aerial photography, surveillance, surveying, and disaster response. However, achieving safe and autonomous navigation in complex and dynamic environments remains a critical research challenge [1], [2].

A widely adopted approach to quadrotor navigation is a two-stage planning framework [3]–[7]. In this framework, a *global* planner first computes a feasible path, which serves as the foundation for collision free polyhedral regions. A *local* planner then dynamically generates smooth and feasible trajectories within these regions, ensuring real-time adaptability while maintaining safety and efficiency. While these approaches perform well in environments with static obstacles, their performance degrades significantly when faced with dynamic obstacles.

Recently, several methods have been proposed to address motion planning with dynamic obstacles [2], [8]–[15]. Despite

Astik Srivastava, Bitla Bhanu Teja, Antony Thomas and Madhava Krishna are with Robotics Research Center, IIT Hyderabad, India. Email:(astik.srivastava, bhanu.teja)@research.iiit.ac.in, (antony.thomas, mkrishna)@iiit.ac.in

Thomas J Chackenkulam is with IIT-BHU, India. Email: thomasj.chackenkulam.cd.mec23@itbhu.ac.in



**Fig. 1:** UAV navigation in a dynamic environment with backup planning. (a) The nominal trajectory (green) becomes unsafe due to dynamic obstacles. (b) A backup trajectory (red) is generated reactively to move the UAV toward a safer region. (c) The UAV tracks the backup trajectory, replanning it as needed to maintain safety. (d) Once a feasible path to the goal becomes available, the UAV switches back to goal-directed planning. Convex collision-free corridors are omitted for clarity. Obstacles within the UAV’s sensing range are shown with red bounding boxes, while those outside the range are shown with green bounding boxes.

this progress, existing approaches remain limited in several key aspects. Many works [8]–[10], [12] consider only a handful of dynamic obstacles, which is far from the complexity encountered in realistic, cluttered environments. Some methods rely on motion capture systems to provide ground-truth obstacle trajectories [2], [9], [13]–[15], an assumption that is impractical for real-world deployment. Furthermore, most approaches [8]–[15] fail to incorporate backup or contingency trajectories. This omission is critical, as in densely populated environments with multiple dynamic obstacles, trajectory optimization alone often fails, leaving the robot without safe fallback (see Fig. 1). Overcoming these limitations is therefore essential for achieving robust and reliable motion planning in dynamic, real-world settings.

In this work, we propose Spatio-Temporal Obstacle-free Trajectory (SPOT), a framework for autonomous navigation

of UAVs in unknown dynamic environments, designed to address the challenges outlined above. Specifically, our key contributions are as follows

- A novel formulation of spatio-temporal RRT\* motion planning for mapless, vision-based navigation in dynamic environments.
- A robust planning pipeline that explicitly handles deadlocks and prevents collisions through backup trajectory generation.
- Extensive validation in both simulation and real-world hardware, demonstrating reliable performance using only onboard sensing and computation. Simulations include scenarios with up to 30 dynamic obstacles. For reproducibility, the code is made available<sup>1</sup>.

## II. RELATED WORK

### A. Global planning and safe flight corridor

A well-established approach to collision avoidance in autonomous navigation involves performing convex decomposition in the configuration space based on the results of a global planner, followed by local trajectory optimization within the decomposed convex corridors [4], [6], [7], [16]–[20]. Convex decomposition is performed around the piecewise-linear path generated by the global planner. Each segment of this path is enclosed within a polyhedron, constructed by first inflating an ellipsoid along the segment and then computing tangent planes at the contact points between the ellipsoid and obstacles. The collection of such overlapping convex polyhedra forms a Safe Flight Corridor (SFC) [17], which defines a collision-free region for trajectory generation.

Global planning algorithms typically employed in this framework include grid-based methods (e.g., Jump Point Search, A\*) [4], [14], [17], as well as sampling-based approaches (e.g., RRT and its variants) [21]. However, state-of-the-art SFC-based methods [4], [7] largely remain restricted to static obstacles in unknown environments, limiting their applicability in dynamic settings.

### B. UAV planning in dynamic environments

A number of recent methods address trajectory planning for UAVs in dynamic environments. One line of work first constructs SFC by considering only static obstacles, and subsequently enforces dynamic collision-avoidance constraints during trajectory optimization. Examples include chance-constrained MPC formulations [8], [10], [11], where dynamic obstacles are typically modeled using simplified geometric primitives such as ellipsoids, spheres, or bounding boxes [1], [8], [22]. An alternative strategy proposed in [23] builds convex hulls around individual dynamic obstacles; however, this approach becomes computationally inefficient as the number of obstacles increases. Lu *et al.* deals with only small moving obstacles such as tennis balls.

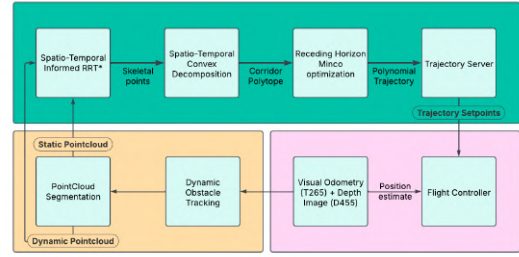


Fig. 2: System architecture of SPOT for reactive UAV motion planning in dynamic, unknown environments.

In contrast, other works model static and dynamic obstacles jointly [2], [13], [15], [24]. Chen *et al.* [24] propose a dual-structure particle-based map to simultaneously capture both obstacle types. Similarly, RAST [13] employs a particle-based occupancy map combined with sampling-based planning to generate trajectories that minimize particle collisions. However, RAST restricts SFC construction to cuboid-shaped corridors, which can be overly conservative in environments with irregular geometries. In [2], static and dynamic obstacles are jointly formulated as collision-avoidance constraints within the trajectory optimization process. More recently, Quan *et al.* [15] utilize a state-time space representation derived from the Euclidean Signed Distance Field (ESDF) to predict pedestrian motion, though their formulation remains limited to 2D environments.

In this work, we construct a spatio-temporal SFC that simultaneously accounts for static and dynamic obstacles. A key distinction is that our framework explicitly incorporates backup trajectories, which are critical in cluttered, dynamic environments where the UAV may otherwise encounter deadlock due to obstacles obstructing its forward motion. Existing approaches [4], [7], [13], [15], [20], [24] do not address this issue. The work of Lu *et al.* [2] is arguably the closest, as it also considers fallback strategies, but their formulation relies on maintaining a static local map. In contrast, our approach (SPOT) is map-less, making it light-weight and more scalable to highly dynamic and unknown environments.

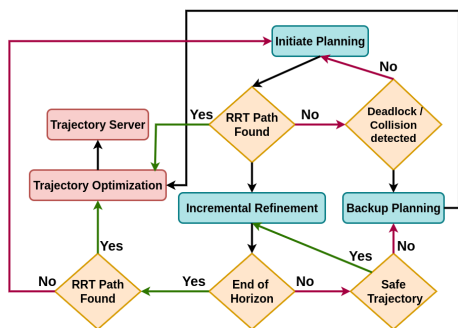
## III. METHODOLOGY

### A. Planning pipeline

In this section, we present our planning architecture for dynamic collision avoidance. An overview of the system is shown in Fig. 2. The proposed framework SPOT consists of three main components:

- **Object Motion Segmentation:** Since planning is performed directly on point clouds (obtained from the onboard depth camera), it is critical to estimate both the position and velocity of points in the scene. To this end, we employ the onboard dynamic object tracking framework [25], which outputs bounding boxes along with object positions and velocities. Each point lying inside a bounding box inherits the corresponding object velocity, while all other points are treated as static. This

<sup>1</sup><https://astik-2002.github.io/ICRA-2026-SPOT/>



**Fig. 3:** FSM governing the spatio-temporal RRT\* planning pipeline, operating in three modes: *Initial*, *Incremental*, and *Backup*.


step produces a spatio-temporal representation of the environment.

- **Collision Free Skeletal Path Generation:** Using the classified point cloud, we generate a guiding path to the goal via spatio-temporal RRT\* (see Section III-B). Unexplored regions outside the sensor’s field of view (FOV) are assumed obstacle-free, consistent with the mapless design. Path generation is regulated by a Finite State Machine (FSM), which serves as the decision-making layer. An overview of the same is given in Fig. 3 and discussed subsequently. The FSM determines whether to initiate a new search, refine an existing path, or invoke the backup planner, thereby ensuring safe and reliable UAV navigation under dynamic conditions.
- **Convex Decomposition and Trajectory Generation:** The skeletal path is then inflated into a safe corridor using convex decomposition, after which a smooth, dynamically feasible trajectory is generated. This ensures collision-free execution while adhering to UAV dynamics.

The planning pipeline is governed by an FSM-based architecture with three modes: *Initial*, *Incremental*, and *Backup*. The initial and incremental planning blocks build upon the framework in [18], with modifications to support efficient spatio-temporal planning. In the initial mode, RRT\* is invoked to plan a path from the UAV’s starting state to the goal. Once a valid path is found, the portion of the tree within the sensing horizon is used for SFC generation and trajectory optimization. In the incremental mode, the path is continuously refined by adding new nodes to the tree while simultaneously evaluating the safety of the trajectory currently being executed by the UAV. Finally, in the backup mode, if the planner fails to find a feasible path within the time budget, or if the UAV is hovering at a potentially unsafe location, a reactive backup planner is activated to generate an immediate collision-free maneuver, thereby reducing the risk of collisions with dynamic obstacles.

### B. Spatio-Temporal RRT

In this section, we present our Spatio-Temporal RRT\* (ST-RRT\*) algorithm for collision avoidance. In a dynamic environment the occupancy of any given location inherently time-dependent. In other words, whether a location is occupied depends not only on the obstacle’s trajectory but also on the



**Fig. 4:** Spatio-temporal RRT path generated in a dynamic environment. The environment provides a point cloud of obstacles, propagated into the future using velocity information. Each sphere represents the safety margin of a node (feasibility set) at its time of arrival. A color gradient from red to blue illustrates the flow of time, from the present (red) toward the future (blue).

arrival time of the robot at that location. To account for this, we augment the standard RRT\* presented in [18] by propagating the robot’s travel time from the *root* node to every node in the tree. Formally, each node is represented as

$$n_i = (x_i, y_i, z_i, t_i)$$

where  $(x_i, y_i, z_i)$  denotes the position in 3D space and  $t_i$  is the accumulated travel time from the root node. This representation allows each node to encode both spatial position and arrival time, thereby enabling collision checks to be performed consistently against obstacle trajectories in the spatio-temporal space. An illustration is provided in Fig. 4.

1) *Travel time estimate:* To incorporate the arrival time at each node, we employ a kinematics-based travel-time estimate, which is propagated and encoded into every node during the ST-RRT\* construction. Let  $v_{max}$  denote the maximum admissible velocity for the robot. Consider a parent node  $n_i = (\mathbf{p}_i, t_i)$ ,  $\mathbf{p}_i = [x_i, y_i, z_i]^T$  and a child node  $n_j = (\mathbf{p}_j, t_j)$ ,  $\mathbf{p}_j = [x_j, y_j, z_j]^T$ . The constructed spatio-temporal tree must satisfy two conditions: (i) causality, that is, the arrival time at any child node must strictly exceed that of its parent, and (ii) velocity feasibility, that is, the child node must be reachable from its parent under the maximum velocity bound  $v_{max}$  (see Fig. 4). Therefore, a child node is admissible if and only if it belongs the feasibility set

$$\mathcal{F}(n_i) = \left\{ n_j \mid t_j > t_i, \|\mathbf{p}_j - \mathbf{p}_i\| \leq v_{max}(t_j - t_i) \right\} \quad (1)$$

This formulation ensures that each edge in the tree is both temporally ordered and kinematically feasible, thereby embedding spatio-temporal consistency directly into the RRT\* expansion process. Moreover, the feasibility set  $\mathcal{F}(n_i)$  naturally extends to incorporate additional kinematic constraints, such as bounded acceleration.

2) *Edge cost:* Given the spatio-temporal nature of the planning problem, we define a heuristic cost function for the connection between a parent node and a child node as

$$C_e = \|\mathbf{p}_j - \mathbf{p}_i\| + w(t_j - t_i) \quad (2)$$

where  $w$  is a weighting parameter which can be interpreted as a *desired velocity*. The above formulation enables a tunable trade-off between spatial distance and temporal separation. This flexibility allows the planner to adapt to application-specific priorities, such as minimizing travel distance or performing cautious motion planning in cluttered dynamic environments. In the steer function for ST-RRT\* and also during rewiring, we ensure that the parent and child nodes are consistent as defined in (1) by restricting the connections of a node to other nodes which are reachable given UAV's velocity constraints.

3) *Sampling strategy*: We have adopt informed sampling<sup>2</sup> to accelerate convergence to an optimal path. Unlike prior work, we extend informed sampling to include not only spatial coordinates but also the time of arrival assigned to each newly added node. This restricts exploration to a spatio-temporal subset of the configuration space guaranteed to contain any better solution.

Formally, let  $c_{\text{best}}$  be the cost of the best path found from start  $\mathbf{p}_s$  to goal  $\mathbf{p}_g$ . The informed set  $\mathcal{X}$  consists of all points  $\mathbf{p}$  satisfying:

$$\|\mathbf{p}_s - \mathbf{p}\| + \|\mathbf{p} - \mathbf{p}_g\| \leq c_{\text{best}} \quad (3)$$

The set  $\mathcal{X}$  forms a prolate hyperspheroid with foci at  $\mathbf{p}_s$  and  $\mathbf{p}_g$ , and major axis length  $c_{\text{best}}$ . Samples are efficiently drawn from this ellipsoid using the closed-form method in [18].

To extend this strategy to spatio-temporal planning, we employ a *decoupled sampling approach*. A candidate spatial sample  $\mathbf{p}_{\text{rand}}$  is first generated within the informed ellipsoid. A time coordinate  $t_{\text{rand}}$  is then assigned to this point by sampling from a distribution designed to prioritize time-efficient solutions

$$t_{\text{rand}} = t_s + k \frac{c_{\text{best}}}{w} \quad (4)$$

where  $k \sim \mathcal{U}(0, 1)$  is uniformly distributed,  $t_s$  is the start time, and  $w$  is the temporal weight from our cost function. Intuitively, this strategy ensures that sampling is focused only on regions of space and time where an improved path could realistically exist. Instead of wasting samples in irrelevant areas or times, the planner concentrates effort within a tightly bounded spatio-temporal subset, thereby accelerating convergence to high-quality, time-efficient trajectories.

4) *Collision checking*: Collision checking in spatio-temporal planning is particularly challenging, as the safety of a node  $n_i$  depends on its distance to obstacles at the specific time  $t_i$  it is visited. For static environments, we use a kd-tree constructed from all the stationary points, which enables efficient nearest-neighbor queries in static point clouds. However, this approach is not directly applicable in dynamic scenes, since obstacle positions and velocities are only known at the instant sensor measurements are received. Beyond that instant, obstacle motion must be propagated forward in time to

<sup>2</sup>Informed RRT\* operates like standard RRT\* until the first solution is found, after which it restricts sampling to a heuristic-defined subset of states, enabling solution refinement while implicitly balancing exploration and exploitation without additional parameters [26].

---

### Algorithm 1 Spatio-Temporal Hash Grid

---

**Input:**  $\mathcal{D}$ , time resolution  $\Delta t$ , cell size  $c$ , start time  $t_{\text{start}}$ , end time  $t_{\text{end}}$

$\triangleright \mathcal{D}$  is defined as the set of dynamic points of the form  $(p, v)$ , where  $p$  denotes the position and  $v$  denotes the associated velocity.

**Output:** Temporal grid  $\mathcal{G}$

```

1:  $\mathcal{G} \leftarrow \emptyset$ 
2: for each  $(p, v) \in \mathcal{D}$  do
3:   for  $t = t_{\text{start}}$  to  $t_{\text{end}}$ ;  $t = t + \Delta t$  do
4:      $p(t) \leftarrow p + v \cdot (t - t_{\text{start}})$ 
5:      $t_{\text{bin}} \leftarrow \lfloor t / \Delta t \rfloor$ 
6:      $(x_i, y_i, z_i) \leftarrow (\lfloor p_x / c \rfloor, \lfloor p_y / c \rfloor, \lfloor p_z / c \rfloor)$ 
7:      $\mathcal{G}[t_{\text{bin}}][(x_i, y_i, z_i)] \leftarrow p(t)$ 

```

---

estimate occupancy at  $t_i$  for safety evaluation. To achieve this efficiently, we employ a spatio-temporal hash grid to perform approximate nearest-neighbor queries. The procedures for hash grid construction and dynamic collision checking are described in Algorithm 1 and Algorithm 2, respectively.

Each dynamic point is encoded in the spatio-temporal hash grid by accounting for the corresponding resolutions in time and space (lines 5–7, Algorithm 1). The resulting hash grid  $\mathcal{G}$  is then queried by Algorithm 1 during collision checking. Specifically, the temporal component of a query point is first used to find the nearest temporal bin of the dynamic obstacle by calling the subroutine `NearestBin` (line 2, Algorithm 2). The spatial component of the query point is then used to construct all 27 neighbors in 3D (lines 3–6, Algorithm 2). Each neighbor is checked for occupancy in  $\mathcal{G}$ , yielding the closest obstacle point (line 11, Algorithm 2).

The safety margin of a node is defined as the minimum of its static clearance (computed from the kd-tree) and its dynamic clearance  $d_{\text{min}}$  (obtained from the spatio-temporal hash grid). A node is considered unsafe if this margin is less than the prescribed safety threshold  $r_{\text{min}}$ .

### C. Spatio-Temporal SFC

As discussed earlier, SFC is defined as a sequence of overlapping convex polyhedra. We use the configuration space convex decomposition presented in [7] for finding spatial corridors. Each polyhedron is constructed around a line segment connecting two points in the free space. Since we employ ST-RRT\* to generate the initial skeletal path, these points correspond to nodes in the path. Each node is associated with a timestamp, enabling spatio-temporal reasoning. Consider two nodes  $n_i$  and  $n_j$ ; conventionally, a polyhedron is constructed around the line segment  $\overline{n_i n_j}$ . This is achieved by inflating an ellipsoid along the segment and computing tangent planes at its contact points with obstacles, which are represented as point clouds.

To extend this formulation to dynamic environments, we discretize the temporal interval between  $n_i$  and  $n_j$  using a step size  $\Delta t$ , resulting in the set  $\mathcal{T}_{ij} = \{t_i, t_i + \Delta t, t_i + 2\Delta t, \dots, t_j\}$ . At each temporal sample  $t \in \mathcal{T}_{ij}$ , the positions

---

**Algorithm 2** Proximity Check to Nearest Obstacle
 

---

**Input:** Query point  $q = (q_x, q_y, q_z, q_t)$ , temporal grid  $\mathcal{G}$ , cell size  $c$ , time resolution  $\Delta t$

**Output:**  $d_{\min}$

▷  $d_{\min}$  represents the minimum distance between the query point and its closest obstacle.

1:  $d_{\min} \leftarrow \infty$ ,  $t \leftarrow q_t$ ,  $p \leftarrow (q_x, q_y, q_z)$

2:  $t_{\text{bin}} \leftarrow \text{NearestBin}(\mathcal{G}, t)$

▷ The subroutine returns the nearest temporal index in  $\mathcal{G}$  corresponding to  $t$ .

3:  $(x_i, y_i, z_i) \leftarrow (\lfloor p_x/c \rfloor, \lfloor p_y/c \rfloor, \lfloor p_z/c \rfloor)$

4: **for**  $dx = -1$  to  $1$  **do**

5:     **for**  $dy = -1$  to  $1$  **do**

6:         **for**  $dz = -1$  to  $1$  **do**

7:              $key \leftarrow (x_i + dx, y_i + dy, z_i + dz)$

8:             **if**  $key \in \mathcal{G}[t_{\text{bin}}]$  **then**

9:                 **for each** obstacle  $o$  in  $\mathcal{G}[t_{\text{bin}}][key]$  **do**

10:                      $d \leftarrow \|p_q - p\|$

11:                      $d_{\min} \leftarrow \min(d_{\min}, d)$

**return**  $d_{\min}$

---

of the dynamic obstacles are predicted under a constant velocity model

$$\mathbf{o}(t) = \mathbf{o}(t_i) + \mathbf{v}_o(t - t_i) \quad (5)$$

where  $\mathbf{o}(t)$  is the predicted obstacle location based on the observed obstacle position  $\mathbf{o}(t_i)$  at time  $t_i$  and  $\mathbf{v}_o$  is the velocity of the obstacle. The corresponding point clouds are aggregated to form the dynamic obstacle set  $\mathcal{O}_{ij}^d$ . The effective obstacle set to be considered during corridor construction is given by


$$\mathcal{O}_{ij} = \mathcal{O}^s \cup \mathcal{O}_{ij}^d \quad (6)$$

where  $\mathcal{O}^s$  denotes the static obstacles. In this manner, the SFC construction inherently accounts for both static and dynamic obstacles across the relevant time horizon. The construction of spatio-temporal collision-free polyhedra and the corresponding evolution of dynamic obstacle point clouds are illustrated in Fig. 5. In our implementation, we empirically fix the temporal resolution to  $\Delta t = 0.2$ s, while constraining the maximum interval  $(t_j - t_i)$  to 2 seconds. These choices strike a balance between prediction fidelity and computational tractability, ensuring that the constant-velocity approximation remains valid over the short prediction horizon.

#### D. Trajectory optimization

Given an initial skeletal path and a sequence of overlapping polyhedra from the SFC, the local planner computes a smooth, dynamically feasible trajectory for a quadrotor. The trajectory  $\mathbf{p}(t) : [0, T] \rightarrow \mathbb{R}^3$  is constrained to lie within the SFC. The total flight time  $T$  is fixed and divided into  $N$  segments with pre-assigned durations  $T_i$ . These durations come from the time assigned to ST-RRT\* during path searching.

The optimization problem is formulated as a nonlinear minimization over the spatial path, leveraging the MINCO framework [5] for efficient computation:



**Fig. 5:** Spatio-temporal convex collision-free polyhedra constructed along the path segment between nodes  $n_i$  and  $n_j$ . The evolution of the safe corridor is color-coded from red to gray over the time span  $[t_i, t_f]$ . During the same interval, the predicted positions of dynamic obstacles are represented as point clouds, also color-coded from  $t_i$  to  $t_f$ . These aggregated point clouds constitute the dynamic obstacle set  $\mathcal{O}_{ij}^d$  used in corridor construction. Final optimized trajectory is visible in green.

$$\min_{\xi} \mathcal{J}_s + \mathcal{J}_d \quad (7a)$$

$$\text{s.t. } \mathbf{p}(t) = \mathcal{M}(\mathbf{q}, \mathbf{T}), \quad \forall t \in T_i, i = 1, \dots, N \quad (7b)$$

$$\mathbf{s}^{(0:2)}(0) = \mathbf{s}_0^{(0:2)}, \quad \mathbf{s}^{(0:2)}(T) = \mathbf{s}_f^{(0:2)} \quad (7c)$$

where  $\mathcal{J}_s$  is the squared-jerk control effort,  $\mathcal{J}_d$  is the penalty function that enforces dynamic limits like maximum velocity, maximum acceleration, etc. and corridor violation penalties.  $\mathcal{M}(\mathbf{q}, \mathbf{T})$  is a linear-complexity smooth map from intermediate points  $\mathbf{q}$  and a time allocation  $\mathbf{T}$  (pre allocated based on ST-RRT\*) for all pieces to the coefficients of polynomial pieces.  $\xi$  denotes the vector whose elements encode the position of  $\mathbf{q}$ , such that each element of  $\mathbf{q}$  can be represented as a convex composition of the vertices of corresponding polygon, with normalized and squared values of  $\xi$  ensuring that all points lie inside the polygon. The constraints in (7c) corresponds to boundary conditions in position, velocity, and acceleration.

The minimum jerk trajectory synthesized by MINCO is parametrized by only  $\mathbf{q}$  and  $\mathbf{T}$ . The gradients of the objective function with respect to  $\xi$  are computed analytically, enabling efficient optimization using the L-BFGS algorithm. This formulation ensures the generated trajectory is safe (contained within the SFC), smooth, and respects the full dynamics of the quadrotor platform.

#### E. Backup trajectory

In this section, we introduce a backup planning algorithm to ensure reactive, collision-free flight when the UAV enters a deadlock or unsafe state. A deadlock is characterized by the failure of local path planning to progress toward the goal,

thereby leaving the UAV exposed to potential collisions with dynamic obstacles.

The central principle of backup planning is to identify an intermediate safe goal in a direction away from obstacles, allowing the planner to reinitiate once the UAV reaches this region. The procedure consists of two stages:

- 1) A convex polygon ( $C_{static}$ ) is constructed around the UAV’s current location, incorporating only static obstacles. This guarantees that the backup maneuver is free from static collisions.
- 2) The repulsive Cauchy Artificial Potential Field (CAPF) approach [27] is used to compute an escape direction vector pointing away from obstacles. The information about the size of the obstacles is taken from the bounding box dimensions received via the object detector. The intersection of this vector with  $C_{static}$  defines the backup goal.

A backup trajectory is then synthesized using ST-RRT\* with convex decomposition and trajectory optimization. The mechanism is triggered when either (i) a collision risk is detected along the current trajectory or (ii) the UAV is hovering while awaiting a feasible path. Once the UAV reaches the backup goal, the finite state machine (FSM) transitions back to its initial mode, thereby resuming its path planning toward the goal.

#### IV. EXPERIMENTS AND VALIDATION

##### A. Simulation results

In this section, we present simulation results for collision avoidance in dynamic, cluttered environments. The results are divided into two categories:

- Simulation with ground truth obstacle information: In this setting, the ground-truth positions and velocities of obstacles within the UAV’s sensing range (6 m) are available, along with a 360° ground-truth point cloud. Simulations are performed in the *gym-pybullet-drones* environment [28] using a ROS 2 interface, and we do not assume perfect tracking conditions, rather relying on realistic simulation of inner loop flight control behaviour. We compare our algorithm against the methods in [14] and [7] under identical conditions.
- Pipeline Validation: To evaluate the complete vision-based navigation pipeline and its potential for sim-to-real transfer, we conduct simulations in the PX4 SITL environment with Gazebo. In this case, obstacle information is obtained from an onboard detector [25], rather than from ground-truth data.

Fig. 6 shows a snapshot from our simulation. All simulations were executed on a workstation equipped with an Intel Core i7 (13th generation) processor and 16 GB of RAM.

1) *Ground-truth simulations*: The simulation environment is a  $16 \times 16$  m<sup>2</sup> area. Dynamic obstacles are modeled as vertical cylinders with a height of 5 m and diameters uniformly sampled between 0.4 m and 1.0 m. Obstacle velocities are randomly sampled within the range 0–0.5 m/s. The maximum

**TABLE I:** Comparison of success rates (%) across methods for environments with increasing obstacle densities.

Method	10 Obstacles	20 Obstacles	30 Obstacles
SPOT	100.0	92	80.2
SPOT-w/o backup	94.2	82.4	52.2
[14]	100.0	71.3	57.1
[7]-5m/s	100.0	70.0	62.2
[7]-1m/s	52.2	42.8	10.2

**TABLE II:** Ablation of SPOT planner across obstacle densities


Obstacles	Length (m)	Replans	Backup	Time (s)
10	16.77	3	0	19.13
20	23.73	11	2	34.56
30	26.44	23	5	50.9

velocity of the UAV is capped at 1 m/s for our algorithm as well as for the method in [14]. For [7], we additionally report results at UAV speeds of 1 m/s and 5 m/s, since the approach is well suited for high-speed navigation. Simulations are conducted in three environments containing 10, 20, and 30 obstacles, respectively. For each environment, 50 independent trials are performed to evaluate the results. In addition, we compare our proposed approach, SPOT, against a variant of SPOT executed without the backup trajectory. The results are summarized in Table I.

It is readily evident from the first two rows of Table I that incorporating a backup trajectory increases the success rate by more than 20%. The improvement in the success rate of [7] at higher UAV velocities can be attributed to the comparatively low speed of obstacles. Specifically, when the UAV moves at 5 m/s, obstacle velocities correspond to at most 10% of the UAV speed, compared to 50% when the UAV moves at 1 m/s. This speed advantage enables more efficient navigation through the environment. However, even under these conditions, [7] performs worse than our approach due to the absence of object motion compensation in dynamic environments.

The primary reason for collisions in [14] is its inability to escape deadlocks, that is, situations where no feasible path to the goal exists. In such cases, the UAV hovers at its current location, making it susceptible to collisions with dynamic obstacles. By contrast, our planner effectively avoids deadlocks by invoking the backup planner, as reflected in the higher success rates reported in Table I. An example of SPOT avoiding a deadlock and switching to a backup trajectory is illustrated in Fig. 1. In Table II we provide mean ablations for the performance parameters of SPOT with different obstacle densities. In the table, Replans and Backup represent the number of times FSM had to switch to *Initial* mode and *Backup* mode respectively.


2) *Pipeline validation*: To validate our sim-to-real pipeline for vision-based dynamic collision avoidance, we use the Gazebo simulation environment with PX4 SITL (see Fig. 7). Humans are simulated as dynamic obstacles, and the integration of our motion planning algorithm with the onboard detector is evaluated. Since the experiments use a depth camera, yaw planning is necessary to ensure the UAV faces



**Fig. 6:** Simulation of UAV avoiding dynamic obstacles. The white polygons represent the spatio-temporal free space around the UAV, which follows the green trajectory. The obstacles within red bounding boxes are with UAVs sensing range, while the obstacles represented as green cylinders lie outside of it.



(a) Sim-to-real validation in the PX4-Gazebo SITL environment for dynamic collision avoidance.




(b) Onboard detector integrated with Spatio-Temporal Path planning

**Fig. 7:** Simulation of UAV avoiding humans with walking/running motion profiles. Bounding Boxes provided by onboard detector [25] are used for segmenting pointclouds, which are then used by Spatio-Temporal planner for generating collision free trajectory

the direction of motion while avoiding dynamic obstacles. For brevity, the yaw planning formulation is not presented here; readers can refer to [29] for details. The simulation is conducted with 5 humans traversing an area of  $20*20\text{ m}^2$ .

### B. Hardware Results

We validate our collision avoidance algorithm through hardware experiments on an in-house developed micro-quadcopter UAV. The platform, as shown in Fig. 8, has a wheelbase of 430 mm and employs 5-inch, three-blade propellers for propulsion. For perception, the UAV is equipped with an Intel RealSense D455, while state estimation is provided by an Intel RealSense T265 camera for visual-inertial odometry. All



**Fig. 8:** Hardware platform.

onboard computation is performed using an NVIDIA Jetson Xavier NX. The UAV runs PX4 on an OmniNXT flight controller to enable trajectory tracking.

We conduct two sets of experiments to evaluate performance in the presence of static and dynamic obstacles, respectively. Unlike [14], we use onboard computation to both track and avoid dynamic obstacles. Representative results are illustrated in Fig. 9.

## V. CONCLUSION

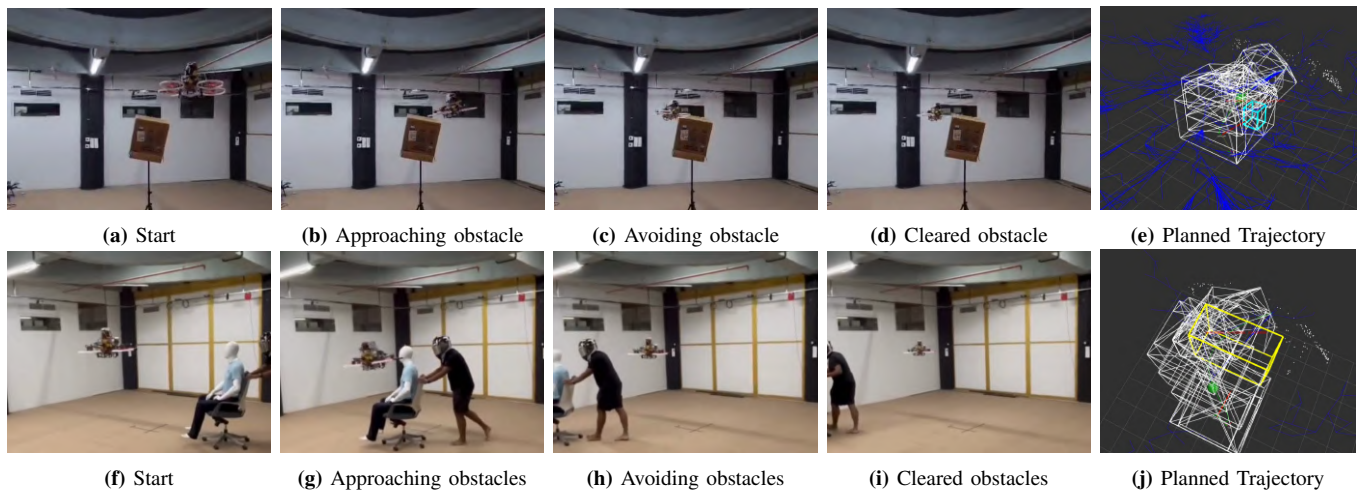
We presented a reactive motion planning framework for quadrotors in unknown, dynamic environments, combining a 4D Spatio-Temporal RRT planner with vision-based Safe Flight Corridor generation and trajectory optimization. Dynamic obstacles are detected and tracked using a vision-based pipeline, and a backup planning module ensures safe navigation in deadlock scenarios. Simulation and hardware experiments demonstrate the method's performance and advantages over state-of-the-art approaches. Future work will focus on improving the estimation of the motion model for dynamic obstacles to enhance prediction and avoidance of collisions. We would also focus on improving the motion planning in more high speed and challenging environments.

## VI. ACKNOWLEDGMENT

The authors acknowledge the support provided by MeitY, Govt. of India, under the project "Capacity Building for Human Resource Development in Unmanned Aircraft System (Drone and Related Technology)."

## REFERENCES

- [1] J. Tordesillas and J. P. How, "Mader: Trajectory planner in multiagent and dynamic environments," *IEEE Transactions on Robotics*, vol. 38, no. 1, pp. 463–476, 2021.
- [2] 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*, vol. 41, pp. 871–886, 2025.
- [3] F. Gao, L. Wang, B. Zhou, X. Zhou, J. Pan, and S. Shen, "Teach-Repeat-Replan: A Complete and Robust System for Aggressive Flight in Complex Environments," *IEEE Transactions on Robotics*, vol. 36, no. 5, pp. 1526–1545, 2020.
- [4] J. Tordesillas, B. T. Lopez, M. Everett, and J. P. How, "Faster: Fast and safe trajectory planner for navigation in unknown environments," *IEEE Transactions on Robotics*, vol. 38, no. 2, pp. 922–938, 2022.
- [5] 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.
- [6] Y. Ren, F. Zhu, W. Liu, Z. Wang, Y. Lin, F. Gao, and F. Zhang, "Bubble planner: Planning high-speed smooth quadrotor trajectories using receding corridors," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 6332–6339.



**Fig. 9:** Hardware experiments demonstrating collision avoidance on an in-house developed UAV. (a–d) UAV avoiding a static obstacle: starting from the initial position, approaching the obstacle, executing avoidance, and clearing it. (e) The planned trajectory. (f–i) UAV avoiding dynamic obstacles: starting from rest, approaching multiple moving obstacles, maneuvering around them, and successfully clearing the scene. (j) The planned trajectory. In the figures, white polygons represent obstacle free space while blue/yellow bboxes represent tracked obstacles

- [7] Y. Ren, F. Zhu, G. Lu, Y. Cai, L. Yin, F. Kong, J. Lin, N. Chen, and F. Zhang, "Safety-assured high-speed navigation for mavs," *Science Robotics*, vol. 10, no. 98, p. ead06187, 2025.
- [8] J. Lin, H. Zhu, and J. Alonso-Mora, "Robust vision-based obstacle avoidance for micro aerial vehicles in dynamic environments," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 2682–2688.
- [9] B. He, H. Li, S. Wu, D. Wang, Z. Zhang, Q. Dong, C. Xu, and F. Gao, "Fast-dynamic-vision: Detection and tracking dynamic objects with event and depth sensing," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 3071–3078.
- [10] Z. Xu, D. Deng, Y. Dong, and K. Shimada, "Dmpc-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.
- [11] J. Hou, X. Zhou, Z. Gan, and F. Gao, "Enhanced decentralized autonomous aerial robot teams with group planning," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 9240–9247, 2022.
- [12] 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.
- [13] G. Chen, S. Wu, M. Shi, W. Dong, H. Zhu, and J. Alonso-Mora, "Rast: Risk-aware spatio-temporal safety corridors for mav navigation in dynamic uncertain environments," *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 808–815, 2023.
- [14] S. Wu, G. Chen, M. Shi, and J. Alonso-Mora, "Decentralized multi-agent trajectory planning in dynamic environments with spatiotemporal occupancy grid maps," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 7208–7214.
- [15] F. Quan, Y. Shen, P. Liu, X. Lyu, and H. Chen, "A state-time space approach for local trajectory replanning of an mav in dynamic indoor environments," *IEEE Robotics and Automation Letters*, 2025.
- [16] R. Deits and R. Tedrake, "Computing large convex regions of obstacle-free space through semidefinite programming," in *Algorithmic Foundations of Robotics XI: Selected Contributions of the Eleventh International Workshop on the Algorithmic Foundations of Robotics*. Springer, 2015, pp. 109–124.
- [17] S. Liu, M. Watterson, K. Mohta, K. Sun, S. Bhattacharya, C. J. Taylor, and V. Kumar, "Planning dynamically feasible trajectories for quadrotors using safe flight corridors in 3-d complex environments," *IEEE Robotics and Automation Letters*, vol. 2, no. 3, pp. 1688–1695, 2017.
- [18] F. Gao, W. Wu, W. Gao, and S. Shen, "Flying on point clouds: Online trajectory generation and autonomous navigation for quadrotors in cluttered environments," *Journal of Field Robotics*, vol. 36, no. 4, pp. 710–733, 2019.
- [19] T. Marcucci, M. Petersen, D. von Wrangel, and R. Tedrake, "Motion planning around obstacles with convex optimization," *Science Robotics*, vol. 8, no. 84, p. eadf7843, 2023.
- [20] Q. Wang, Z. Wang, M. Wang, J. Ji, Z. Han, T. Wu, R. Jin, Y. Gao, C. Xu, and F. Gao, "Fast iterative region inflation for computing large 2-d/3-d convex regions of obstacle-free space," *arXiv preprint arXiv:2403.02977*, 2024.
- [21] C. Richter, A. Bry, and N. Roy, "Polynomial trajectory planning for aggressive quadrotor flight in dense indoor environments," in *Robotics Research: The 16th International Symposium ISRR*. Springer, 2016, pp. 649–666.
- [22] M. Kamel, J. Alonso-Mora, R. Siegwart, and J. Nieto, "Robust collision avoidance for multiple micro aerial vehicles using nonlinear model predictive control," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 236–243.
- [23] J. Tordesillas and J. P. How, "Panther: Perception-aware trajectory planner in dynamic environments," *IEEE Access*, vol. 10, pp. 22 662–22 677, 2022.
- [24] G. Chen, P. Peng, P. Zhang, and W. Dong, "Risk-aware trajectory sampling for quadrotor obstacle avoidance in dynamic environments," *IEEE Transactions on Industrial Electronics*, vol. 70, no. 12, pp. 12 606–12 615, 2023.
- [25] Z. Xu, X. Zhan, Y. Xiu, C. Suzuki, and K. Shimada, "Onboard dynamic-object detection and tracking for autonomous robot navigation with rgb-d camera," *IEEE Robotics and Automation Letters*, vol. 9, no. 1, pp. 651–658, 2024.
- [26] J. D. Gammell, S. S. Srinivasa, and T. D. Barfoot, "Informed rrt\*: Optimal sampling-based path planning focused via direct sampling of an admissible ellipsoidal heuristic," in *2014 IEEE/RSJ international conference on intelligent robots and systems*. IEEE, 2014, pp. 2997–3004.
- [27] A. Srivastava, V. Vasudevan, R. N. Harikesh, and P. Sujit, "A modified artificial potential field for uav collision avoidance," in *Proc. Int. Conf. Unmanned Aircr. Syst.(ICUAS)*, 2023, pp. 499–506.
- [28] J. Panerati, H. Zheng, S. Zhou, J. Xu, A. Prorok, and A. P. Schoellig, "Learning to fly—a gym environment with pybullet physics for reinforcement learning of multi-agent quadcopter control," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 7512–7519.
- [29] B. Zhou, J. Pan, F. Gao, and S. Shen, "Raptor: Robust and perception-aware trajectory replanning for quadrotor fast flight," *IEEE Transactions on Robotics*, vol. 37, no. 6, pp. 1992–2009, 2021.