

# PUMA: Fully Decentralized Uncertainty-aware Multiagent Trajectory Planner with Real-time Image Segmentation-based Frame Alignment

Kota Kondo, Claudius T. Tewari, Mason B. Peterson, Annika Thomas,  
 Jouko Kinnari, Andrea Tagliabue, Jonathan P. How

**Abstract**—Fully decentralized, multiagent trajectory planners enable complex tasks like search and rescue or package delivery by ensuring safe navigation in unknown environments. However, deconflicting trajectories with other agents and ensuring collision-free paths in a fully decentralized setting is complicated by dynamic elements and localization uncertainty. To this end, this paper presents (1) an uncertainty-aware multiagent trajectory planner and (2) an image segmentation-based frame alignment pipeline. The uncertainty-aware planner propagates uncertainty associated with the future motion of detected obstacles, and by incorporating this propagated uncertainty into optimization constraints, the planner effectively navigates around obstacles. Unlike conventional methods that emphasize explicit obstacle tracking, our approach integrates implicit tracking. Moreover, sharing trajectories between agents can cause potential collisions due to frame misalignment. Addressing this, we introduce a novel frame alignment pipeline that rectifies inter-agent frame misalignment. This method leverages a zero-shot image segmentation model for detecting objects in the environment and a data association framework based on geometric consistency for map alignment. Our approach accurately aligns frames with only 0.18 m and 2.7° of mean frame alignment error in our most challenging simulation scenario. In addition, we conducted hardware experiments and successfully achieved 0.29 m and 2.59° of frame alignment error. Together with the alignment framework, our planner ensures safe navigation in unknown environments and collision avoidance in decentralized settings.

## SUPPLEMENTARY MATERIAL

**Video:** <https://youtu.be/W73p42XRcaQ>

**Code:** <https://github.com/mit-acl/puma>


## I. INTRODUCTION

Multiagent Unmanned Aerial Vehicle (UAV) trajectory planning has been extensively studied [1]–[16]. In real-world deployments of multiagent trajectory planning methods, it is crucial to (1) detect and avoid collisions with **obstacles**, (2) correct for **localization errors and uncertainties**, and (3) achieve **scalability** to a large number of agents.

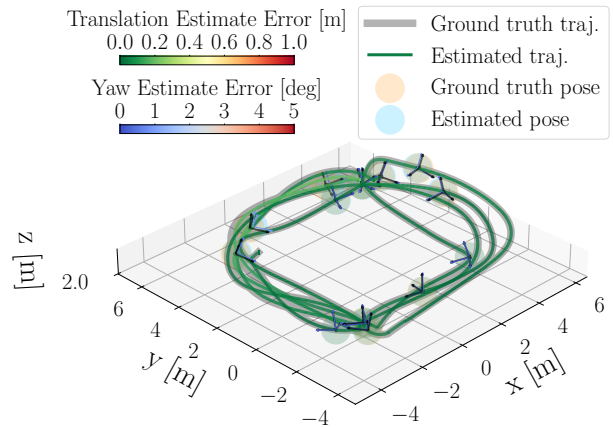
Equipping agents with sensors, often cameras, is a common method to detect and avoid unfamiliar obstacles [17]–[24]. This provides agents with real-time situational awareness, facilitating informed decisions for collision avoidance in dynamic settings. However, these sensors typically have a restricted field of view (FOV), and therefore the orientation

The authors are with the Department of Aeronautics and Astronautics, Massachusetts Institute of Technology. {kkondo, cttewari, masonbp, annikat, jkinnari, atagliab, jhow}@mit.edu.

This work is supported by Boeing Research & Technology and the Air Force Office of Scientific Research MURI FA9550-19-1-0386.



(a) Multiagent Frame Alignment Evaluation Hardware Environment.



(b) Tracking Quality in Case 30 (difficulty: **Moderate**): Our pipeline successfully estimates the drifted state of vehicle 1 even in multiagent environments. The estimated trajectory demonstrates low errors, indicated in green, and closely aligns with the ground truth trajectory. The estimated poses also exhibit a high degree of overlap with the ground truth, as shown by the significant overlapping of circles. The coordinate frames are captured at intervals of 5 s.

Fig. 1. PUMA is thoroughly evaluated in simulation and hardware experiments are also performed to evaluate the real-time image segmentation-based frame alignment pipeline.

of the UAV must be considered when planning trajectories through new spaces, and hence planners with limited FOV must be perception-aware to reduce environment uncertainty.

In addition, while many existing planners operate under the assumption of perfect global localization, often relying on Global Navigation Satellite Systems (GNSS), for practical real-world applications where perfect global localization is not available, agents need to perform onboard localization, such as visual-inertial odometry (VIO). However, onboard localization in environments without global information introduces the significant challenge of overcoming state estimation drift. Because estimation drift corrupts an agent’s estimate of its own location, any other spatial information

that is considered relative to the agent’s location (e.g. future trajectories and locations of observed obstacles) will be corrupted as well. So for a decentralized planning pipeline to function, it is important to both estimate the true relative poses between agents and to understand the uncertainty associated with that estimate. If these quantities are accurately represented, agents can effectively communicate about their trajectories and environment.

In multiagent trajectory planning, there are two approaches: centralized and decentralized. Centralized planners [8], [9] have one entity that directs all agent trajectories, providing simplicity but with the drawback of a potential single point of failure. Conversely, decentralized planners [4]–[7], [12], [14] enable each agent to plan its own trajectory, offering better scalability and resilience. Similarly synchronous planners [7], [11], [25] require synchronization before planning, whereas asynchronous methods [4]–[6] let agents plan individually, proving to be more scalable.

To tackle the aforementioned challenges: (1) **unknown object detection and collision avoidance**, (2) **localization errors and uncertainties**, and (3) **scalability**, we introduce **PUMA**. Our contributions include:

- 1) A fully decentralized, asynchronous, Perception-aware, and Uncertainty-aware MultiAgent trajectory planner.
- 2) An inter-agent frame alignment pipeline that utilizes an image segmentation technique and a data association approach using geometric consistency.
- 3) Comprehensive simulation benchmarks, comparing PUMA with a state-of-the-art perception-aware planner in **100** simulations, and thorough evaluation of our segmentation-based frame alignment pipeline across **42** different scenarios.
- 4) Hardware tests on UAVs, demonstrating real-time on-board frame alignment, ensuring successful frame synchronization.

## II. UNCERTAINTY-AWARE MULTIAGENT TRAJECTORY PLANNING

This section describes our uncertainty-aware planner. Our planner generates collision-free position and yaw trajectories in a decentralized manner while simultaneously accounting for various uncertainties in the environment. Specifically, it tackles uncertainties from (1) the prediction of previously detected objects and agents, (2) imperfect vision-based detection, and (3) potential unknown obstacles in an agent’s path. The notation used is provided in Table I.

### A. Uncertainty Propagation for Known Obstacles/Agents

Predicting accurate future trajectories of dynamic obstacles from their observation history and accounting for localization errors of other agents is challenging. Therefore, we need to factor in the uncertainty of these trajectories for robust collision avoidance.

Kamel et al. [26] employed the Extended Kalman Filter (EKF) for uncertainty propagation of dynamic objects’ trajectories:

$$P_{k+1} = F_k P_k F_k^T \quad (1)$$

TABLE I. Notation

Symbol	Definition
$(\cdot)_k$	Variable at time step $k$ .
$P$	Covariance matrix.
$H$	Observation matrix.
$F$	State-transition matrix defined as $e^{At}$ , where $A$ is a time-invariant system dynamic matrix.
$\hat{P}_{k k-1}$	Estimate of $P$ at time $k$ given observations up to and including time $k-1$ .
$S$	Innovation covariance.
$K$	Optimal Kalman gain.
$\mathbf{x}; \mathbf{x}^o$	Ego agent’s state; obstacle/peer-agents’ state.
$R_{\max}$	Maximum covariance of the observation noise.
$\theta$	Opening angle of the cone that approximates FOV.
$f_{FOV}$	Function that determines if a given point $\mathbf{p}$ is in FOV. Defined as $-\cos(\theta/2) + (\mathbf{p})_z / \ \mathbf{p}\ $ .
$R_{FOV}$	Covariance of observation noise that dynamically changes uncertainty propagation. Defined as $R_{\max} \frac{1-f_{FOV}}{1+f_{FOV}+\epsilon}$ , where $\epsilon$ is a small number.
$(\cdot)^m$	Variable in the direction of motion uncertainty formulation.
$\mathbf{p}_{\text{traj}}$	Points uniformly sampled from the ego agent’s trajectory.
$\mathbf{v}_{\text{limit}}$	Upper bound of velocity constraint (constant).
$P_{\text{pos}}$	Position covariance matrix.
$\mathbf{v}_{\max}(P_{\text{pos}}^m)$	Velocity constraint (dependent of $P_{\text{pos}}^m$ ).
$\mathcal{F}_w$	Global world frame.
$\mathcal{F}_{\ell_i}$	Local and potentially drifted frame of the $i$ -th agent—if an agent has perfect localization $\mathcal{F}_{\ell_i} = \mathcal{F}_w$ .
$\mathcal{F}_i$	$i$ -th agent’s body frame.
$\mathcal{F}_c$	Camera frame.
$X_\alpha^\beta$	Pose ( $\in SE(d)$ where $d$ is 2 or 3) of $\alpha$ in the frame $\mathcal{F}_\beta$ or the rigid body transformation from frame $\mathcal{F}_\alpha$ to frame $\mathcal{F}_\beta$ .
$\mathbf{p}_{\text{cent}}^\alpha$	Position of a centroid in the frame $\mathcal{F}_\alpha$
$u_{\text{cent}}, v_{\text{cent}}$	Centroid camera pixel coordinates.
$K_c$	Camera intrinsic calibration matrix ( $\in \mathbb{R}^{3 \times 3}$ ).
$\lambda$	Positive scalar ( $\in \mathbb{R}^+$ ).
$\kappa$	Number of timesteps to keep a landmark without new measurements before deleting.
$\delta_{i,o}$	Time since the $i$ -th agent last observed landmark $o$ .

The 95% confidence interval error ellipsoid from  $P$  is inflated as a minimum allowable acceptable distance between the ego agent and obstacle—a larger  $P$  leads to a bigger boundary. While effective, this method does not include perception information. We propose a modification to the propagation equation using the prediction and measurement update steps of a linear Kalman Filter:

$$\begin{aligned} \hat{P}_{k|k-1} &= F_k P_{k-1|k-1} F_k^T \\ S_k &= H_k \hat{P}_{k|k-1} H_k^T + R_{k,FOV}(\mathbf{x}_k, \mathbf{x}_k^o) \\ K_k &= \hat{P}_{k|k-1} H_k^T S_k^{-1} \\ P_{k|k} &= (I - K_k H_k) \hat{P}_{k|k-1} \end{aligned} \quad (2)$$

where the state of each obstacle contains its  $\mathbb{R}^3$  position, velocity, and acceleration. If obstacles or other agents are within FOV, the predicted uncertainty is reduced, simulating the process of obtaining new measurements on the position of the peer agents and obstacles.

### B. Uncertainty Propagation to Safely Navigate in Unknown Space

When navigating in unknown environments, the agent should monitor for unknown obstacles along the direction in which the ego agent is moving (unknown space). This is obtained by introducing an additional uncertainty propagation, where we modify the innovation covariance formulation

in Eq. (2) as:

$$S_k^m = H_k^m \hat{P}_{k|k-1}^m H_k^{mT} + R_{k,FOV}^m(\mathbf{x}_k, \mathbf{p}_{\text{traj}}), \quad (3)$$

where  $R_{k,FOV}^m(\mathbf{x}_k, \mathbf{p}_{\text{traj}})$  takes the ego agent's state  $\mathbf{x}_k$  and  $\mathbf{p}_{\text{traj}}$ , a set of points uniformly sampled from the ego agent's future trajectory. The propagated uncertainty remains minimal if the points in  $\mathbf{p}_{\text{traj}}$  are within FOV. As  $P_{pos}^m$  enlarges, indicating increasing uncertainty along the direction of motion, we impose tighter constraints on the planner's maximum velocity.

$$\mathbf{v}_{\max,k}(P_{pos,k}^m) = \frac{\bar{v}_{\text{limit}}}{\sqrt{\text{diag}(P_{pos,k}^m)}} \quad (4)$$

Note that this operation is done element-wise for the  $x$ ,  $y$ , and  $z$  components.

### C. Uncertainty-aware Optimization Formulation

This section presents the uncertainty-aware optimization formulation, which leverages the future propagation of  $P$  and optimizes the vehicle trajectory over a time horizon, similar to an MPC formulation. The optimization problem is formulated as follows:

$$\begin{aligned} \min_{\substack{\mathbf{p}(t), \psi(t), \\ \mathbf{n}_{ij}, d_{ij}, T}} \quad & \alpha_j \int_{t_{in}}^T \|\dot{\mathbf{j}}\|^2 dt + \alpha_\psi \int_{t_{in}}^T (\dot{\psi})^2 dt + \alpha_T T \\ & + \alpha_{\mathbf{g}_P} \|\mathbf{p}(T) - \mathbf{g}_P\|^2 + \alpha_{g_\psi} \|\psi(T) - g_\psi\|^2 \\ \text{s.t.} \quad & \mathbf{x}(t_{in}) = \mathbf{x}_{in}, \quad \mathbf{v}(T) = \mathbf{0}, \\ & \mathbf{a}(T) = \mathbf{0}, \quad \dot{\psi}(T) = 0, \\ & \mathbf{n}_{ij}^T(P_{pos})\mathbf{q} + d_{ij}(P_{pos}) < 0 \quad \forall \mathbf{q} \in \mathcal{Q}_j \forall i, j \\ & \text{abs}(\mathbf{v}) \leq \mathbf{v}_{\max}(P_{pos}^m) \quad \forall \mathbf{v} \in \mathcal{V}_j, \forall j \\ & \text{abs}(\mathbf{a}_l) \leq \mathbf{a}_{\max} \quad \forall l \in L_P \setminus \{n_P - 1, n_P\} \\ & \text{abs}(\dot{\mathbf{j}}_l) \leq \dot{\mathbf{j}}_{\max} \quad \forall l \in L_P \setminus \{n_P - 2, n_P - 1, n_P\} \\ & \text{abs}(\psi_l) \leq \psi_{\max} \quad \forall l \in L_\Psi \setminus \{n_\Psi\} \end{aligned}$$

Equations (2) and (3)

Our formulation adopted the notation from PANTHER [22]. It is essential to note that the variables for the separating plane ( $\mathbf{n}_{ij}$  and  $d_{ij}$ ) depend on  $P_{pos}$ , and the velocity constraint is dependent of  $P_{pos}^m$ . Importantly, this formulation does not necessitate explicit obstacle tracking. If a trajectory that leads to obstacles being in FOV of the agent provides a smoother and shorter path, the planner is naturally incentivized to track the obstacles, similarly for the direction of motion. Our approach is the first to demonstrate implicit obstacle tracking in this way.

### D. Multiagent Trajectory Deconfliction

To share the planned trajectories with other agents, we use our Robust MADER trajectory deconfliction framework [5] to guarantee safe and collision-free trajectories. This framework, by introducing (1) a Delay Check step, (2) a two-step trajectory publication scheme, and (3) a trajectory storing-and-checking approach, gives agents the ability to deconflict trajectories asynchronously and without centralized communication.

## III. FRAME ALIGNMENT PIPELINE

To address discrepancies in agents' local frames caused by onboard localization drift, we introduce a pipeline inspired by [27] where drifted frames of each agent are aligned in real-time and using only distributed communication. The goal of frame alignment is to find the rigid transformation  $X_{\ell_j}^{\ell_i}$  that maps poses from agent  $j$ 's local frame,  $\mathcal{F}_{\ell_j}$ , into agent  $i$ 's local frame,  $\mathcal{F}_{\ell_i}$ . To estimate  $X_{\ell_j}^{\ell_i}$ , each agent  $i$  creates maps of static landmarks in the environment and then aligns its map with each of its neighbors  $j$ . Maps are created using observations that leverage recent image segmentation techniques [28], [29]. This image segmentation allows agents to create maps of previously unencountered objects, removing the need for objects of known classes to exist in the environment. This workflow is illustrated in Fig. 2.

Map creation begins with detecting segments from an image frame using a generic zero-shot segmentation model [29] and extracting the pixel coordinates ( $u_{\text{cent}}, v_{\text{cent}}$ ) segments' centroids. Segments are filtered based on size, and then each centroid position is computed with

$$p_{\text{cent}}^{\ell_i} = \mathbf{X}_c^{\ell_i} (\lambda K_c^{-1} [u_{\text{cent}} \quad v_{\text{cent}} \quad 1]^T). \quad (5)$$

We make the assumption that landmarks are on the ground plane, which constrains the  $z$  element of  $p_{\text{cent}}^{\ell_i}$  to be 0, further defining the value of  $\lambda$  and consequently, the  $x$  and  $y$  components of  $p_{\text{cent}}^{\ell_i}$ . In practice, observed objects will often be 2.5D (i.e., resting on the ground plane but protruding from the plane), so measurement covariances are stretched along the range direction to represent the uncertainty accurately.

Once object centroids and their associated measurement covariances have been obtained in the drifted world frame, we create maps and perform frame alignment. Identified centroids are associated with existing map objects using the global nearest neighbor approach [30], which formulates a linear assignment problem based on the Mahalanobis distance between new centroid measurements and existing map objects. New landmarks are created in the map when object measurements cannot be associated with existing landmarks, and landmarks are deleted from the map after  $\kappa$  frames of not being seen to eliminate any drift distortion in the map.

Subsequently, pairs of distributed robots align their drifted frames using their landmarks maps. Before maps can be aligned, pairs of objects within each map must be associated together. This is accomplished using CLIPPER [31] which performs global data association between objects using geometric consistency. Once associations have been made, a weight  $W(\delta_{i,o}, \delta_{j,o})$  is applied to each pair of associated objects such that  $W(\delta_{i,o}, \delta_{j,o}) = (\delta_{i,o} \delta_{j,o})^{-1}$ , which prioritizes using landmarks with recent observations. Finally, the transformation  $X_{\ell_j}^{\ell_i}$  is computed using a weighted version of Arun's method [32], with weights computed using  $W(\delta_{i,o}, \delta_{j,o})$ . This transformation is applied to trajectories that an agent receives from its neighbor  $j$  to rectify frame misalignment.




Fig. 2. Frame Alignment Pipeline Workflow: The process starts by un-distorting a raw, fisheye image. Next, the pipeline identifies and segments individual objects within the image. We then determine each identified object’s centroid, filtering out objects based on size. Once processed, these centroids are projected onto a 2D map, which is used in the frame alignment among agents.

## IV. SIMULATION RESULTS

### A. Single-agent Uncertainty-aware Planner Benchmarking

This section evaluates the performance of our uncertainty-aware planner through simulations. We compared our planner with PANTHER\* [33], a state-of-the-art perception-aware planner that has an *explicit obstacle tracking term* in its cost function. We tested both planners in a single-agent setting, where the agent is tasked to fly through one obstacle that follows a pre-determined trefoil trajectory. Fig. 3 shows the flight environment, and Table II summarizes the results of 100 flight simulations. The metrics used to evaluate the performance are as follows:

- 1) Travel time: duration to complete the path.
- 2) Computation time: duration to replan at each step.
- 3) Number of collisions: the number of collisions that the agent experiences.
- 4) Known obstacle FOV rate: the percentage of time that the agent keeps known obstacles within its FOV.
- 5) Known obstacle continuous FOV detection: Consecutive time that an obstacle is continuously kept within the FOV of the agent.
- 6) Unknown space FOV rate: the percentage of time that the agent keeps unknown space (direction of motion) within its FOV.
- 7) Unknown space continuous FOV detection: consecutive time that unknown space (direction of motion) is continuously kept within the FOV of the agent.

TABLE II. Single Agent Uncertainty-aware Planner Benchmarking — 100 simulations

Planner	Travel Time [s]	Comp. Time [ms]	# Colls.	Known Obst.		Unknown Space	
				FOV Rate [%]	Conti. FOV Dets. [s]	FOV Rate [%]	Conti. FOV Dets. [s]
PUMA (proposed)	5.2	5712	0	62.8	2.4	13.0	0.59
PANTHER* [33]	4.5	1891	0	100	4.5	0.1	0.0

Table II highlights a key distinction between PUMA and PANTHER\* [33]. Specifically, PUMA places emphasis on perceiving unknown spaces while concurrently maintaining awareness of known obstacles. PUMA retains perception of the known obstacles for 62.8% of the flight time while focusing on the unknown space (direction of motion) for 13.0%. In contrast, PANTHER\*, due to the *explicit obstacle*

*tracking term* in its cost function, remains focused on known obstacles for 100% of the flight duration, giving no attention to unknown spaces, shown by a rate of 0.1%.

Fig. 3 illustrates PUMA’s capability to balance uncertainties about obstacles and unknown spaces. Specifically, at  $t=1.0$  s, as depicted in Fig.3b, the agent focuses on the known obstacle, decreasing its uncertainty while uncertainty about unknown space goes up. Conversely, at  $t=2.5$  s, the agent’s attention shifts to unknown spaces it approaches, reducing its uncertainty as indicated in Fig. 3e. This dynamic management of uncertainties allows the agent to navigate unknown environments safely.

While PUMA takes more computation time than PANTHER\*, the imitation learning approach used for DeepPANTHER [33], can be used in future work to significantly reduce the computation time of PUMA while maintaining very similar levels of performance.

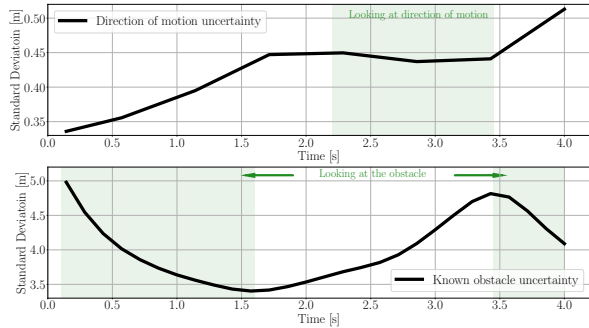
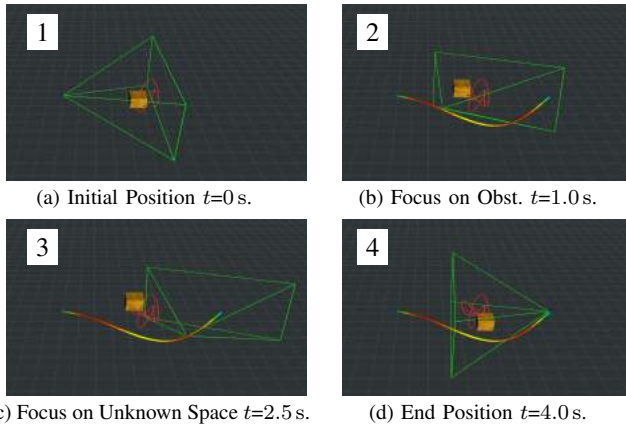
### B. Frame Alignment Pipeline Evaluation

This section evaluates our image segmentation-based frame-alignment pipeline. We use simulated environments involving two flying agents and artificially introduce either constant bias or linear drift in translation and yaw to the state estimation of one agent. Our pipeline effectively detects these misalignments and estimates the relative states.

1) *Simulation Environments*: We tested the pipeline using two sets of objects: flat pads and random objects. As mentioned in Section III, our pipeline assumes that objects are on the ground. Consequently, it performs better with flat pads depicted in Fig. 4a. However, to further evaluate its robustness, we also examined it with various randomly chosen objects, as illustrated in Fig.4b.

For the simulations, we employed the camera model from Intel® RealSense™ T265. The raw images, as displayed in Figs. 4a and 4b, were fed directly into our segmentation-based pipeline.

2) *Simulation Results*: We evaluate our pipeline by considering two different trajectories: (1) agents following the same circular trajectory and (2) partially overlapping circular (POC) trajectory. For each trajectory, we tested three different drift scenarios: (1) no drift, (2) constant frame offset, and (3) linearly added drift. We introduced these drifts and offsets to vehicle 1 in a) translation only, b) yaw only, and



(e) Uncertainty Growth: When solving optimization problem, the planner propagates uncertainty of obstacles and unknown space.

Fig. 3. PUMA balances reducing uncertainties of (1) known obstacles and (2) potential obstacles in trajectory: The green pyramid represents the FOV. The agent navigates around the dynamic obstacle to reach its destination on the opposite side. The trajectory color shows the agent’s velocity —red is fast, and blue slow.




Fig. 4. Pipeline Evaluation Simulation Environments.

c) both translation and yaw. When applied, the constant bias consists of 1 m along the  $x$  and  $y$  axes and 10 deg in yaw, and linear drift accumulates at a rate of 0.05 m/s along  $x$  and  $y$  and 0.05 deg/s in yaw. Accurately aligning frames is more difficult when drift accumulates linearly than when there is only a constant bias, and drift in both translation and yaw presents a more difficult scenario than drift in only translation or only yaw.

We first showcase the frame alignment results of our pipeline under artificially introduced constant drifts in Fig. 5. The pipeline first identifies the pads in the environment and places them into a 2D map, and it performs frame alignment in real-time and estimates vehicle 1’s drifted state. The estimated state and ground truth values in Fig. 5 show that the pipeline successfully estimates the drifted state of vehicle

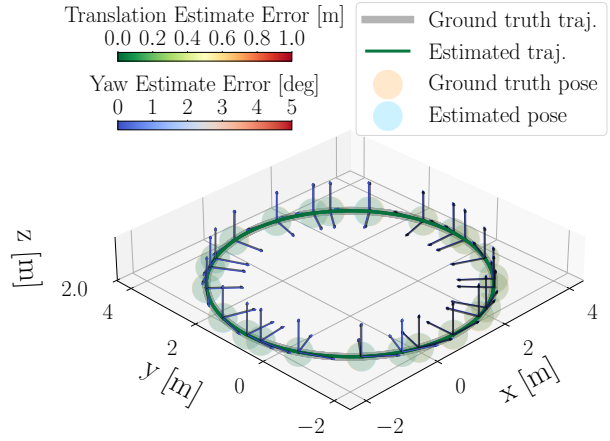


Fig. 5. Frame alignment quality in Case 3 (difficulty: **Easy**): Trajectory color shows the estimation error associated with it and the color of the coordinate frame attached to vehicle 1 shows the yaw estimation error. The pipeline successfully estimates the drifted state of vehicle 1 with minimal errors (Table III). The coordinate frames are captured at intervals of 5 s.

TABLE III. Simulation Benchmarking 1 — Object Type: Pads

Environment Settings					Results					
Case	Difficulty	Traj.	Added Error Type	Translation/Yaw	X Err. [m] Mean	X Err. [m] Std.	Y Err. [m] Mean	Y Err. [m] Std.	Yaw Err. [deg] Mean	Yaw Err. [deg] Std.
1	Very Easy	Circle	None	None	0.0	0.02	0.0	0.02	0.01	0.1
2	Easy	POC	None	None	0.09	0.05	0.07	0.07	0.14	0.23
3	Easy	Const.	Const.	Translation	0.0	0.02	0.0	0.02	0.02	0.09
4	Easy	Circle	Const.	Yaw	0.0	0.01	0.0	0.02	0.05	0.1
5	Moderate	Const.	Const.	Both	0.0	0.02	0.0	0.02	0.05	0.13
6	Easy	Const.	Const.	Translation	0.13	0.05	0.03	0.03	0.14	0.28
7	Easy	POC	Const.	Yaw	0.09	0.07	0.06	0.1	0.75	0.97
8	Moderate	Const.	Const.	Both	0.09	0.07	0.02	0.08	1.14	1.52
9	Moderate	Circle	Linear	Translation	0.01	0.11	0.02	0.11	2.42	0.14
10	Moderate	Circle	Linear	Yaw	0.0	0.01	0.0	0.01	0.52	0.12
11	Hard	Circle	Linear	Both	0.01	0.1	0.03	0.1	1.92	0.14
12	Hard	POC	Linear	Translation	0.13	0.11	0.05	0.1	2.7	0.22
13	Hard	POC	Linear	Yaw	0.11	0.02	0.02	0.01	0.54	0.23
14	Very Hard	POC	Linear	Both	0.14	0.09	0.04	0.09	2.22	0.19

1 with minimal errors. Tables III and Table IV summarize the results of our simulations. The tables are organized as follows: (1) Columns 1–4 show the trajectory and drift types, (2) Columns 5–7 show the drift values, (3) the last 6 columns show the mean and standard deviation of the  $x$ ,  $y$ , and yaw errors, respectively. Table III VI shows that our pipeline successfully estimates the drifted state of vehicle 1 in all cases. It is apparent that the error increases for the more difficult cases, but even in the most challenging case, the mean yaw error is less than  $2.5^\circ$ , and the mean position error is less than 0.21 m. Note that the random objects are 3D objects and, thus, are not flat on the ground. Therefore, it is challenging to detect them; however, our pipeline still successfully estimates the drifted state of vehicle 1.

### C. Multiagent Uncertainty-Aware Planner Evaluation on Frame Alignment Pipeline

This section tests the uncertainty-aware planner presented in Section II, integrated with the frame alignment pipeline discussed in Section III. Following the same pattern as in Section III, both constant and linearly increasing drifts were introduced in environments with pads and random objects. A distinguishing feature here, as opposed to Section IV-B, is that the agents’ trajectories are not simple circles or POC trajectories. Instead, they are shaped by the uncertainty-aware

TABLE IV. Simulation Benchmarking 2 — Object Type: Random

Environment Settings					Results					
Case	Difficulty	Traj.	Added Offset/ Drift Type	Translation/ Yaw	X Err. [m]		Y Err. [m]		Yaw Err. [deg]	
					Mean	Std.	Mean	Std.	Mean	Std.
15	Easy	Circle	None	None	0.01	0.02	0.01	0.02	0.16	0.3
16	Moderate	POC	None	None	0.05	0.0	0.03	0.0	0.76	0.0
17	Moderate	Circle	Const.	Translation	0.0	0.02	0.0	0.02	0.09	0.26
18	Moderate		Const.	Yaw	0.0	0.02	0.01	0.02	0.1	0.26
19	Hard		Const.	Both	0.0	0.03	0.01	0.02	0.05	0.34
20	Moderate	POC	Const.	Translation	0.1	0.08	0.04	0.08	0.34	0.77
21	Moderate		Const.	Yaw	0.16	0.11	0.07	0.09	1.1	2.3
22	Hard		Const.	Both	0.14	0.06	0.04	0.04	0.84	0.69
23	Hard	Circle	Linear	Translation	0.01	0.11	0.03	0.12	2.49	0.28
24	Hard		Linear	Yaw	0.01	0.02	0.0	0.03	0.63	0.24
25	Very Hard		Linear	Both	0.0	0.1	0.02	0.1	2.02	0.25
26	Very Hard	POC	Linear	Translation	0.11	0.24	0.01	0.13	1.4	1.21
27	Very Hard		Linear	Yaw	0.05	0.3	0.08	0.21	1.51	1.05
28	Very Hard		Linear	Both	0.12	0.12	0.02	0.16	1.68	1.16

TABLE V. Multiagent Simulation Benchmarking — PUMA on Frame Alignment Pipeline on Pads

Environment Settings					Results						
Case	Difficulty	Added Offset/ Drift Type	Translation/ Yaw		X Err. [m]		Y Err. [m]		Yaw Err. [deg]		# Colls.
					Mean	Std.	Mean	Std.	Mean	Std.	
29	Moderate	None	None		0.0	0.04	0.0	0.03	0.15	0.45	0
30	Moderate	Const.	Translation		0.02	0.06	0.01	0.04	0.04	0.47	0
31	Moderate	Const.	Yaw		0.02	0.04	0.01	0.03	0.25	0.34	0
32	Hard	Const.	Both		0.02	0.03	0.0	0.03	0.11	0.28	0
33	Very Hard	Linear	Translation		0.02	0.2	0.04	0.23	0.01	1.78	0
34	Very Hard	Linear	Yaw		0.01	0.07	0.02	0.06	0.7	0.38	0
35	Very Hard	Linear	Both		0.18	0.19	0.05	0.17	1.31	0.88	0

TABLE VI. Multiagent Simulation Benchmarking — PUMA on Frame Alignment Pipeline on Random Objects

Environment Settings					Results						
Case	Difficulty	Added Offset/ Drift Type	Translation/ Yaw		X Err. [m]		Y Err. [m]		Yaw Err. [deg]		# Colls.
					Mean	Std.	Mean	Std.	Mean	Std.	
36	Hard	None	None		0.0	0.05	0.02	0.06	0.38	0.43	0
37	Hard	Const.	Translation		0.06	0.06	0.04	0.05	0.05	0.76	0
38	Hard	Const.	Yaw		0.05	0.06	0.05	0.05	0.11	0.74	0
39	Very Hard	Const.	Both		0.02	0.07	0.02	0.06	0.01	0.87	0
40	Very Hard	Linear	Translation		0.17	0.19	0.05	0.16	0.03	1.56	0
41	Very Hard	Linear	Yaw		0.01	0.12	0.03	0.13	1.47	0.72	0
42	Very Hard	Linear	Both		0.02	0.37	0.02	0.17	0.47	2.15	0

planner, introducing an extra layer of complexity to the frame alignment task. Yet, despite this additional complexity in the vehicle motion, our pipeline successfully estimates frame misalignment. As Tables V and VI indicate, the pipeline reliably approximates the drifted state of vehicle 1 in every scenario. Figures 1b and 6 further depicts how our image segmentation-based method estimate drifts in a multiagent environment.

## V. HARDWARE EXPERIMENTS FRAME ALIGNMENT PIPELINE EVALUATION

To first evaluate our sparse mapping capability, we conducted an experiment where a single agent flew at a speed of 0.5 m/s for 60 s, tracing a 2.5 m-circle. The UAV in use was equipped with both an Intel® RealSense™ T265 and an Intel® NUC™ 10. Note that our image segmentation-based frame alignment pipeline operated in real-time on the NUC™ 10, powered by the Intel® Core™ i7-10710U Processor. By employing FastSAM [29], we produce segmented images approximately every 200 ms. We arranged flat mats on the ground for this test, as visualized in Fig. 1a.

To assess our sparse mapping and segmentation-based pipeline, we utilized ground truth localization from a motion capture system. This way, we can assess the sparse mapping capability by matching the generated data with the ground truth locations of the objects. It also allows us to keep track of the ground truth frame discrepancy, which is zero, and, therefore, we can assess the pipeline accurately.

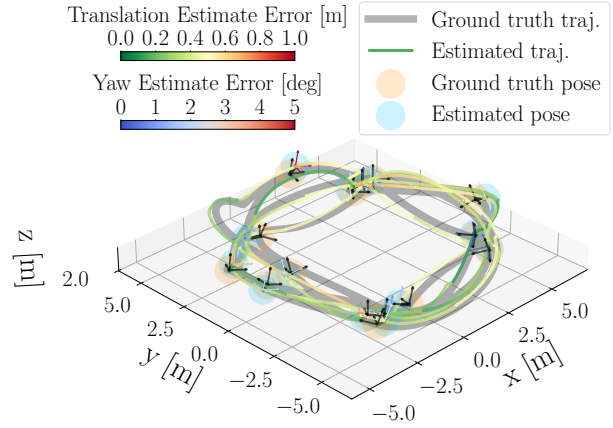


Fig. 6. Tracking Quality in Case 42 (difficulty: **Very Hard**): Even in the hardest case [linearly increasing drift in x, y, and yaw, random objects, PUMA trajectory], our pipeline successfully estimates the drifted state of vehicle 1. The coordinate frames are captured at intervals of 5 s.

The result, as illustrated in Fig. 7, confirms that our image segmentation-based mapping technique effectively extracted a sparse map directly from raw images.

In addition, we performed an assessment of the effectiveness of our frame alignment pipeline in a multiagent context by having two agents navigate a 2.5 m-circle (see Fig. 1a). Table VII shows that our frame alignment method adeptly synchronized the coordinate frames of the two agents in real time.




Fig. 7. Single-agent hardware evaluation: agent follows 2.5-m-circle while successfully creating sparse map.

TABLE VII. Multiagent Frame Alignment Hardware Evaluation

Settings				Results				
Traj.	Added Offset/ Drift Type	Translation/ Yaw		X Err. [m]		Y Err. [m]		Yaw Err. [deg]
				Mean	Std.	Mean	Std.	Mean
Circle	None	None		0.29	0.36	0.16	0.26	2.59
								1.77

## VI. CONCLUSIONS

In multiagent trajectory planning, the challenges of frame alignment drift and obstacle and peer collision avoidance must all be addressed. This paper introduced a perception and uncertainty-aware planner that can navigate unknown spaces while avoiding known obstacles. Our real-time, image segmentation-based pipeline robustly estimates frame alignment between drifted agent frames. Future work includes larger-scale flight experiments and extending our pipeline to 3D mapping.

## REFERENCES

- [1] G. Ryou, E. Tal, and S. Karaman, "Cooperative Multi-Agent Trajectory Generation with Modular Bayesian Optimization," in *Robotics: Science and Systems XVIII*. Robotics: Science and Systems Foundation, Jun. 2022.
- [2] P. Peng, W. Dong, G. Chen, and X. Zhu, "Obstacle avoidance of resilient UAV swarm formation with active sensing system in the dense environment," *arXiv preprint arXiv:2202.13381*, 2022.
- [3] Y. Gao, Y. Wang, X. Zhong, T. Yang, M. Wang, Z. Xu, Y. Wang, Y. Lin, C. Xu, and F. Gao, "Meeting-merging-mission: A multi-robot coordinate framework for large-scale communication-limited exploration," in *2022 IEEE/RSJ IROS*, 2022, pp. 13 700–13 707.
- [4] J. Tordesillas and J. P. How, "MADER: Trajectory planner in multi-agent and dynamic environments," *T-RO*, 2021.
- [5] K. Kondo, R. Figueroa, J. Rached, J. Tordesillas, P. C. Lusk, and J. P. How, "Robust MADER: Decentralized multiagent trajectory planner robust to communication delay in dynamic environments," *arXiv preprint arXiv:2303.06222*, 2023.
- [6] X. Zhou, J. Zhu, H. Zhou, C. Xu, and F. Gao, "EGO-Swarm: A Fully Autonomous and Decentralized Quadrotor Swarm System in Cluttered Environments," 2020.
- [7] B. Sabetghadam, R. Cunha, and A. Pascoal, "A distributed algorithm for real-time multi-drone collision-free trajectory replanning," *Sensors*, vol. 22, no. 5, 2022.
- [8] D. R. Robinson, R. T. Mar, K. Estabridis, and G. Hower, "An Efficient Algorithm for Optimal Trajectory Generation for Heterogeneous Multi-Agent Systems in Non-Convex Environments," *IEEE RA-L*, vol. 3, no. 2, pp. 1215–1222, Apr. 2018.
- [9] J. Park, J. Kim, I. Jang, and H. J. Kim, "Efficient Multi-Agent Trajectory Planning with Feasibility Guarantee using Relative Bernstein Polynomial," in *ICRA*, May 2020, pp. 434–440.
- [10] J. Hou, X. Zhou, Z. Gan, and F. Gao, "Enhanced decentralized autonomous aerial swarm with group planning," *ArXiv*, vol. abs/2203.01069, 2022.
- [11] R. Firoozi, L. Ferranti, X. Zhang, S. Nejadnik, and F. Borrelli, "A distributed multi-robot coordination algorithm for navigation in tight environments," *arXiv preprint arXiv:2006.11492*, 2020.
- [12] C. Toumeh, "Decentralized multi-agent planning for multirotors: a fully online and communication latency robust approach," *arXiv preprint arXiv:2304.09462*, 2023.
- [13] Z. Wang, C. Xu, and F. Gao, "Robust trajectory planning for spatial-temporal multi-drone coordination in large scenes," in *2022 IEEE/RSJ IROS*, 2022, pp. 12 182–12 188.
- [14] S. Batra, Z. Huang, A. Petrenko, T. Kumar, A. Molchanov, and G. S. Sukhatme, "Decentralized control of quadrotor swarms with end-to-end deep reinforcement learning," in *Conference on Robot Learning*. PMLR, 2022, pp. 576–586.
- [15] B. Şenbaşlar and G. S. Sukhatme, "Dream: Decentralized real-time asynchronous probabilistic trajectory planning for collision-free multi-robot navigation in cluttered environments," *arXiv preprint arXiv:2307.15887*, 2023.
- [16] B. Şenbaşlar, P. Luiz, W. Hönig, and G. S. Sukhatme, "Mrnav: Multi-robot aware planning and control stack for collision and deadlock-free navigation in cluttered environments," *arXiv preprint arXiv:2308.13499*, 2023.
- [17] J. Thomas, J. Welde, G. Loianno, K. Daniilidis, and V. Kumar, "Autonomous flight for detection, localization, and tracking of moving targets with a small quadrotor," *IEEE RA-L*, vol. 2, no. 3, pp. 1762–1769, 2017.
- [18] B. Penin, R. Spica, P. R. Giordano, and F. Chaumette, "Vision-based minimum-time trajectory generation for a quadrotor uav," in *2017 IEEE/RSJ IROS*, 2017, pp. 6199–6206.
- [19] D. Falanga, P. Foehn, P. Lu, and D. Scaramuzza, "PAMPC: Perception-aware model predictive control for quadrotors," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 1–8.
- [20] V. Murali, I. Spasojevic, W. Guerra, and S. Karaman, "Perception-aware trajectory generation for aggressive quadrotor flight using differential flatness," in *2019 American Control Conference (ACC)*, 2019, pp. 3936–3943.
- [21] I. Spasojevic, V. Murali, and S. Karaman, "Perception-aware time optimal path parameterization for quadrotors," in *2020 ICRA*, 2020, pp. 3213–3219.
- [22] J. Tordesillas and J. P. How, "Panther: Perception-aware trajectory planner in dynamic environments," *IEEE Access*, vol. 10, pp. 22 662–22 677, 2022.
- [23] B. Zhou, J. Pan, F. Gao, and S. Shen, "Raptor: Robust and perception-aware trajectory replanning for quadrotor fast flight," *T-RO*, vol. 37, no. 6, pp. 1992–2009, 2021.
- [24] X. Wu, S. Chen, K. Sreenath, and M. W. Mueller, "Perception-aware receding horizon trajectory planning for multicopters with visual-inertial odometry," *IEEE Access*, vol. 10, pp. 87 911–87 922, 2022.
- [25] R. Van Parys and G. Pipeleers, "Distributed model predictive formation control with inter-vehicle collision avoidance," in *ASCC*. IEEE, 2017.
- [26] 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)*, 2017, pp. 236–243.
- [27] M. B. Peterson, P. C. Lusk, and J. P. How, "MOTLEE: Distributed mobile multi-object tracking with localization error elimination," in *IEEE/RSJ IROS*, 2023, pp. 719–726.
- [28] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, P. Dollár, and R. Girshick, "Segment anything," *arXiv:2304.02643*, 2023.
- [29] X. Zhao, W. Ding, Y. An, Y. Du, T. Yu, M. Li, M. Tang, and J. Wang, "Fast segment anything," 2023.
- [30] Y. Bar-Shalom and X.-R. Li, *Multitarget-multisensor tracking: principles and techniques*. YBs Storrs, CT, 1995, vol. 19.
- [31] P. C. Lusk, K. Fathian, and J. P. How, "CLIPPER: A graph-theoretic framework for robust data association," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 13 828–13 834.
- [32] K. S. Arun, T. S. Huang, and S. D. Blostein, "Least-squares fitting of two 3-d point sets," *IEEE TPAMI*, no. 5, pp. 698–700, 1987.
- [33] J. Tordesillas and J. P. How, "Deep-PANTHER: Learning-based perception-aware trajectory planner in dynamic environments," *IEEE RA-L*, vol. 8, no. 3, pp. 1399–1406, 2023.