

# PRIMER: Perception-Aware Robust Learning-based Multiagent Trajectory Planner

Kota Kondo<sup>1</sup>, Claudius T. Tewari<sup>1</sup>, Andrea Tagliabue<sup>1</sup>, Jesus Tordesillas<sup>2</sup>,  
Parker C. Lusk<sup>1</sup>, Mason B. Peterson<sup>1</sup>, Jonathan P. How<sup>1</sup>

**Abstract**—In decentralized multiagent trajectory planners, agents need to communicate and exchange their positions to generate collision-free trajectories. However, due to localization errors/uncertainties, trajectory deconfliction can fail even if trajectories are perfectly shared between agents. To address this issue, we first present PARM and PARM\*, perception-aware, decentralized, asynchronous multiagent trajectory planners that enable a team of agents to navigate uncertain environments while deconflicting trajectories and avoiding obstacles using perception information. PARM\* differs from PARM as it is less conservative, using more computation to find closer-to-optimal solutions. While these methods achieve state-of-the-art performance, they suffer from high computational costs as they need to solve large optimization problems onboard, making it difficult for agents to replan at high rates. To overcome this challenge, we present our second key contribution, PRIMER, a learning-based planner trained with imitation learning (IL) using PARM\* as the expert demonstrator. PRIMER leverages the low computational requirements at deployment of neural networks and achieves a computation speed up to 5614 times faster than optimization-based approaches.

## I. INTRODUCTION

In recent years, multiagent UAV trajectory planning has been extensively studied [1]–[22]. In real-world deployments of multiagent trajectory planning methods, it is crucial to deal with challenges such as (1) detecting and avoiding collisions with **unknown obstacles**, (2) handling **localization errors/uncertainties**, (3) achieving **scalability** to a large number of agents, and (4) enabling **fast and efficient computation** for onboard replanning and quick adaptation to dynamic environments. However, finding effective solutions to these challenges remains an open question.

One approach to address challenges such as detecting and avoiding unknown obstacles, even in the presence of localization errors and uncertainties, is to equip each agent with a sensor, typically a camera, to perceive the surrounding environment. This allows agents to gather real-time information about their surroundings, enabling them to make informed decisions and take appropriate actions to avoid collisions and navigate through dynamic environments. However, this sensor often has a limited field of view (FOV), making the orientation of the UAV crucial when planning trajectories

<sup>1</sup>The authors are with the Department of Aeronautics and Astronautics, Massachusetts Institute of Technology. {kkondo, cttewari, atagliab, plusk, kkondo, jhow}@mit.edu.

<sup>2</sup>The author is with the Institute for Research in Technology, ICAI School of Engineering, Comillas Pontifical University (Spain). jtordesillas@comillas.edu

This work is supported by Boeing Research & Technology and AFOSR MURI FA9550-19-1-0386

TABLE I. State-of-the-art UAV Trajectory Planners

Method	Multiagent	Perception-aware	Learning-based
EGO-Swarm [7]	Yes	No	No
DMPC [23]			
MADER [4]			
decMPC [24]			
Micro Swarm [18]			
RMADER [5], [6]			
MRNAV [16]			
DREAM [19]			
HDSM [17]			
Raptor [25]			
Time-opt [26]			
PANTHER [27]			
PA-RHP [28]			
PA-RHP [29]			
Deep-PANTHER [21]	No	Yes	Yes
CGD [30]			
LPA [31]			
PUMA [20]			
PRIMER (proposed)	Yes	Yes	No
PRIMER (proposed)	Yes	Yes	Yes

through unknown space. Therefore, planners for flying with limited FOV sensors generally need to be perception-aware to ensure that as many obstacles or other UAVs as possible are kept within the FOV.

When scaling multiagent trajectory planners, it is important to note that, with centralized planners, each agent needs to listen to a single entity that plans all the trajectories [9], [10]. While this approach simplifies planning, the central entity may act as a single point of failure, and the replanning abilities of the agent depend on their ability to communicate with the central entity. Decentralized planners greatly mitigate these issues, as each agent plans its own trajectory [4], [5], [7], [8], [13], [15]. Decentralized planners are therefore generally considered to be inherently more scalable and robust to failures.

Similarly, synchronous planners such as [8], [12], [22] require all agents to wait at a synchronization barrier until planning can be globally triggered, whereas asynchronous planning enables each agent to independently trigger the planning step without considering the planning status of other agents. Asynchronous methods are typically more scalable compared to synchronous methods [4], [5], [7].

Many optimization-based approaches [4], [5], [7], [13], [20] have been proposed for multiagent trajectory generation. However, these methods often require significant computational resources, making them challenging to deploy in dynamic environments where fast, on-the-fly replanning is essential. In particular, when optimizing perception-aware trajectories while considering multiagent interactions, com-

putation times increase significantly, as demonstrated in Section III and [20], which hinders real-time deployment. To address this issue, researchers have explored imitation learning (IL)-based approaches [21], [32], [33], which offer the advantage of faster replanning while still producing near-optimal trajectory generation.

To tackle the challenges of (1) **unknown objects detection and collision avoidance**, (2) **localization errors/uncertainties**, (3) **scalability**, and (4) **fast and efficient computation**, we propose PRIMER, an IL-based decentralized, asynchronous, perception-aware multiagent trajectory planner. Table I provides a comparison of PRIMER with state-of-the-art approaches, and our contributions include:

- 1) PARM/PARM\* – decentralized, asynchronous, perception-aware, multiagent trajectory planner.
- 2) PRIMER – IL-based decentralized, asynchronous, perception-aware, multiagent approach for translational and yaw trajectory generation with the use of Long Short-Term Memory (LSTM) neural network.
- 3) Extensive simulation benchmarking, we compared the performance of PARM, PARM\*, and PRIMER.

## II. PERCEPTION-AWARE MULTIAGENT TRAJECTORY GENERATION

This section outlines our perception-aware multiagent planning approach, encompassing both optimization-based and IL-based methods. First, we introduce PARM and PARM\*, which are optimization-based perception-aware multiagent trajectory planners. These planners utilize the Robust MADER trajectory deconfliction framework [6] to ensure safe and collision-free trajectories. By considering both position and yaw, PARM and PARM\* are capable of tracking multiple obstacles while generating trajectories. Next, we present PRIMER, an IL-based planner that is trained using PARM\* as an expert. PRIMER offers fast computation by leveraging the benefits of IL. By incorporating the expertise of PARM\*, PRIMER is able to generate effective trajectories while significantly reducing computational requirements. As Table II shows PRIMER is the first learning-based perception-aware multiagent trajectory planner designed to track multiple obstacles while accounting for both position and yaw.

### A. PARM / PARM\* —Optimization-based Decentralized, Asynchronous Perception-Aware Multiagent Planning

Our previous work, PANTHER [27], proposed a perception-aware trajectory planner for a single agent in dynamic environments, generating trajectories to avoid obstacles while keeping them in the sensor FOV. In [21], PANTHER\* improved the original PANTHER with less conservatism by including separating planes and trajectory total time as optimization variables, but both were limited to tracking and avoiding only one obstacle at a time. To address this limitation, we modify the optimization problem solved by PANTHER to enable the tracking and avoidance of multiple obstacles, resulting in PARM—a decentralized,

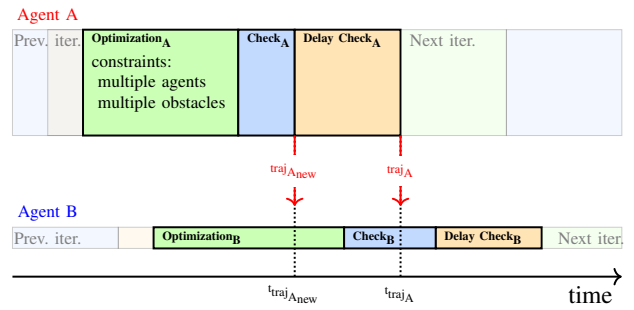


Fig. 1. PARM’s trajectory optimization and deconfliction sequence: PARM uses an optimization-based approach to generate trajectories for each agent, followed by a conflict detection and resolution step based on the Robust MADER framework. Each agent first generates a new trajectory in the optimization step, and then checks if there are any conflicts with the trajectories received from other agents. If no conflicts are detected, the agent publishes its new trajectory and begins checking for potential collisions in a delay check step. This delay check step is a sequence of checks over a period of time. Finally, if no conflicts are detected during the delay check, the agent commits to the new trajectory and publishes it. However, if conflicts are detected, the agent reverts to the trajectory from the previous iteration and discards the new trajectory. More details on the RMADER approach can be found in Section II of [6].

asynchronous, perception-aware multi-agent trajectory planning system. PARM incorporates this modified optimization approach into the RMADER [6] deconfliction framework, which solves a feasibility linear problem between its own committed trajectories and the shared trajectories from other agents (see Fig. 1 and [4], [6] for details). To facilitate the tracking of multiple obstacles and agents, we modify the FOV term as described in Section IV of [27] as follows.

$$-\alpha_{FOV} \sum_i^n \left[ \int_0^T \{\text{inFOV}(\text{obstacle}_i)\}^3 dt \right] \quad (1)$$

where  $\alpha_{FOV}$  is the weight for the FOV term,  $n$  is the number of obstacles,  $T$  is the total time of the trajectory,  $\text{inFOV}()$  returns a higher number when  $\text{obstacle}_i$  is in FOV.

Additionally, following the method outlined in [21], which minimizes trajectory time in the optimization formulation and consequently reduces trajectory costs, we introduce PARM\*, a method that generates closer-to-optimal trajectories. However, this more complex formulation increases computation time, resulting in slower replanning rates.

Note that coupled trajectory generation has the advantage of enabling the optimizer to simultaneously consider both position and yaw trajectories, taking into account the impact of each on the other. On the other hand, the decoupled approach involves optimizing either the position or yaw trajectory first, and then optimizing the other trajectory based on the pre-determined trajectory. While this approach may reduce computation and complexity, it could also result in sub-optimal position and yaw trajectories overall, as the two trajectories are not jointly optimized. Therefore we designed PARM\* in a way that they optimize both position and yaw in a coupled manner (See Table II).

Note that we assume our system is differential flat [34], and the flat outputs are position, yaw, and their derivatives.

## B. PRIMER —Imitation Learning-based PARM\*

In our previous work, Deep-PANTHER [21], we used IL to train a neural network that generates a desired position trajectory, while relying on closed-form solutions to obtain the direction where the onboard sensor should be looking at (e.g., yaw on a multirotor). This closed-form yaw solution generates yaw trajectories given position trajectories, reducing the output dimension of the learned policy. However, this approach is not scalable in multi-obstacle environments since the closed-form solution only generates yaw trajectories for a single given obstacle. To address this limitation, we design PRIMER using a multi-layer perceptron (MLP) that generates both position and yaw trajectories. To generate both position and yaw trajectories, PRIMER has the size of the neural network to 4 fully connected layers, each with 1024 neurons, and is trained to imitate the optimal perception-aware trajectories generated by PARM\*.

Additionally, we added a Long Short-Term Memory (LSTM) [35] feature-extraction network to the MLP used in PRIMER, inspired by the ground-robot motion planning approach [36]. This allowed the neural network to accept various numbers of obstacles and agents as input, whereas traditional fully-connected feedforward neural networks can only handle a fixed number of obstacles. LSTM can take as many obstacles and agents as possible and generate a fixed length of the latent output, which we feed into the fully connected layers.

It is also worth noting that IL-based approaches are more scalable in practice than optimization-based approaches. As the number of agents and obstacles in the environment increases, optimization-based approaches need to include more constraints in the optimization, leading to significant computational requirements. On the other hand, IL-based approaches such as PRIMER are able to handle larger-scale environments with little to no additional computational overhead with the use of LSTM.

Fig. 2 shows an architecture of PRIMER. We first feed the predicted trajectories of obstacles and received other agents’ trajectories to the LSTM, which outputs a fixed size vector. We then combine it with the agent’s own state and feed this into the fully connected layers.

Table II shows the comparison of the state-of-the-art perception-aware trajectory planners. PARM/PARM\* are the first perception-aware multiagent trajectory planner that generates position and yaw coupled trajectory while tracking multiple obstacles, and PRIMER is the extension of PARM, in the sense that it has much faster computation time, leveraging IL-based planner.

Similar to PARM/PARM\*, PRIMER leverages Robust MADER’s trajectory deconfliction approach, ensuring that all trajectories generated by PRIMER undergo deconfliction checks to guarantee collision-free trajectory generation between agents. Unlike traditional optimization-based methods, which ensure collision-free trajectories by incorporating trajectories from other agents into their constraints, NN-based planners do not inherently provide such guarantees. However,

the overall architecture ensures that the final committed trajectories at the conclusion of the replanning steps are collision-free.

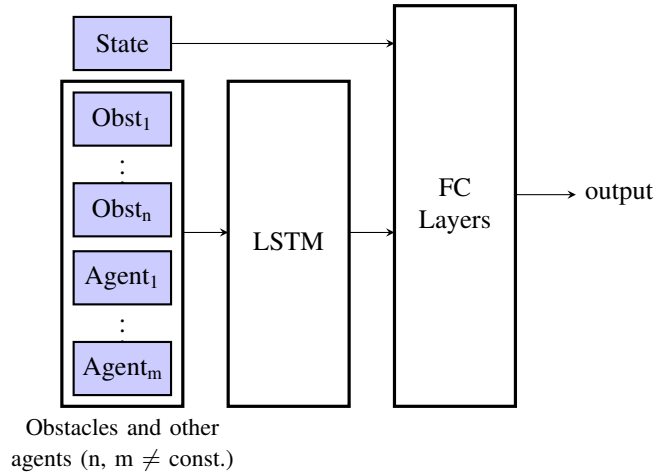


Fig. 2. PRIMER Network Architectures

TABLE II. State-of-the-art Perception-aware Obstacle Tracking Trajectory Planners

Method	Tracking Multi-obstacles	Multi-agents	Trajectory	Planning
[37]	No	No	Only Position	Optimization-based (slow & not scalable)
[38]	No	No	Position & Yaw	Optimization-based (slow & not scalable)
PANTHER / PANTHER* [21], [27]	No	No	Position & Yaw	Optimization-based (slow & not scalable)
Deep-PANTHER [21]	No	No	Only Position <sup>1</sup>	IL-based (faster & scalable)
PARM / PARM* (proposed)	Yes	Yes	Position & Yaw	Optimization-based (slow & not scalable)
PRIMER (proposed)	Yes	Yes	Position & Yaw	IL-based (faster & scalable)

## C. PRIMER Training Setup

We used the student-expert IL learning framework, where PARM\* acts as an expert that provides demonstrations, and PRIMER is the student, trained so that its neural network can reproduce the provided demonstrations. We trained the student in an environment that contained multiple dynamic obstacles flying a randomized trefoil-knot trajectory, as well as other PRIMER agents. The terminal goal for the student was also randomized. To collect the data and train the student, we utilized the Dataset-Aggregation algorithm (Dagger) [39], and Adam [40] optimizer. Additionally, we introduced a weighted loss function between position and yaw. During the training process, we found that it was more difficult to train the yaw trajectory than the position trajectory, and thus we weighted the yaw loss function. In our training, we set the weight  $\alpha$  to 70. The total loss is defined as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{pos}} + \alpha \mathcal{L}_{\text{yaw}} \quad (2)$$

<sup>1</sup>Deep-PANTHER [21] generates only position trajectory, and yaw trajectory is generated by closed-form solution based on the position trajectory.

where  $\mathcal{L}_{\text{total}}$  is the total loss,  $\mathcal{L}_{\text{pos}}$  is the loss for the position,  $\alpha$  is the weight between the position and yaw trajectories, and  $\mathcal{L}_{\text{yaw}}$  is the loss for the yaw trajectory.

#### D. Obstacle Sharing

As shown in Fig. 3, each agent detects and tracks obstacles and shares their predicted trajectories with other agents. This obstacle-sharing architecture allows the agents to have a better understanding of the surrounding environment as a team.

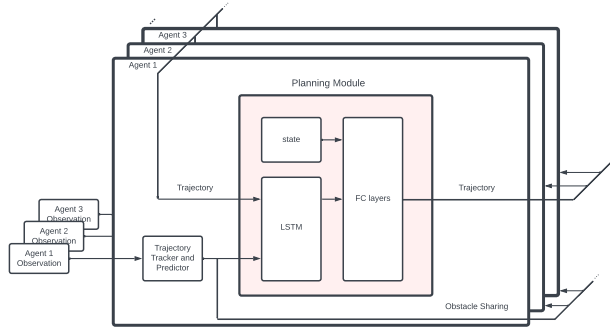


Fig. 3. PRIMER Planning and Sharing Trajectory Architecture

### III. SIMULATION RESULTS

#### A. PRIMER vs. PARM\* in single-agent, single-obstacle environment

Table III compares the average performance of PRIMER and PARM\* in a simulation environment with a single dynamic obstacle following a trefoil trajectory, while the agent flies diagonally to avoid obstacles. The comparison is based on the average trajectory cost, computation time, obstacle avoidance failure rate, and dynamic constraint failure rate. As shown in Table III, PRIMER achieves a **5614-time** reduction in computation time, with only a 6.7% increase in cost, while maintaining a 100% success rate and 0% dynamic constraint violations.

TABLE III. PRIMER vs. PARM\* in single-agent, single-obstacle environment

	Avg. Cost	Compu. Time [ms]	Success Rate [%]	Obst. Avoidance, Dyn. Constr. Failure Rate [%]
PARM*	403	4495.8	100	0.0
PRIMER	431	0.8008	100	0.0

#### B. Multiagent and multi-obstacle benchmarking

We also tested PARM, PARM\*, and PRIMER in two different environments: (1) one agent and two obstacles and (2) two agents and two obstacles. To conduct the experiment, we positioned the agents in a 3.0 m radius circle and had them exchange positions diagonally, as shown in Fig. 4. The obstacles are boxes with 0.5-m edges. We set the maximum dynamic limits to 2.0 m/s, 10.0 m/s<sup>2</sup>, and 30.0 m/s<sup>3</sup> for velocity, acceleration, and jerk, respectively.

We conducted all simulations on an Alienware Aurora r8 desktop running Ubuntu 20.04, which is equipped with an

Intel® Core™ i9-9900K CPU clocked at 3.60 GHz with 16 cores and 62.6 GiB of RAM.

Table IV compares the average performance of PARM, PARM\*, and PRIMER with (1) one agent with two obstacles, and (2) three agents with two obstacles. The metrics used to evaluate the performance are as follows:

- 1) Computation time: the time to replan at each step.
- 2) Success rate: the rate at the agents successfully reach the goal without collisions.
- 3) Travel time: the time it takes for the agent to complete the position exchange.
- 4) FOV rate: the percentage of time that the agent keeps obstacles within its FOV when the agent is closer than its camera's depth range.
- 5) Number of continuous FOV detection frames: the number of consecutive frames that an obstacle is kept within the FOV of the agent.
- 6) Translational dynamic constraints violation rate: the violation rate of the max velocity, acceleration, and jerk.
- 7) Yaw dynamic constraints violation rate: the violation rate of the maxi yaw rate.

Note that for PARM and PARM\*, we had them generate both 1 and 6 trajectories per replanning and compared their performance. When an agent generates 6 trajectories, it selects the one with the lowest cost. Although 1-trajectory replanning is faster, 6-trajectory replanning is more likely to find a better trajectory. Table IV shows the 1-trajectory replanning approach is significantly faster. However, in more challenging environments with three agents, PARM\*'s 1-trajectory approach has a lower success rate compared to the 6-trajectory replanning approach.

Overall, all three methods achieve successful position exchange with similar performance. However, PRIMER outperforms PARM and PARM\* significantly in terms of computation time.

In a more complex environment with three agents and two obstacles, both PARM and PRIMER achieve a high success rate, while PARM\* falls short of achieving a 100% success rate. This is because agents spend too much time optimizing their trajectories in PARM\*, causing the optimization constraints to become outdated and leading to conflicts during the check and delay check steps. PARM\* also suffers from long computation time, resulting in very few generated trajectories passing the check and delay check steps, thus leading to a low success rate.

Fig. 5 illustrates (a) Computation Time, (b) Travel Time, and (c-d) Trajectory Smoothness. While the 6-trajectory approach finds smoother trajectories, it requires much longer computation time, resulting in increased travel time. The trajectories generated by PRIMER exhibit smoothness comparable to those of PARM\*, while PRIMER have significantly faster computation times compared to PARM and PARM\*.

### IV. CONCLUSIONS

In conclusion, our work has addressed the critical issue of trajectory deconfliction in perception-aware, decentralized multiagent planning. We first presented PARM and PARM\*,

TABLE IV. Simulation Benchmarking

Env.	Method	# Trajs	Avg. Compu. Time [ms]	Success Rate [%]	Avg. Travel Time [s]	FOV Rate [%]	Avg. of Max # Conti. FOV Detection Frames	Translational Dyn. Constraints Violation Rate [%]	Yaw Dyn. Constraints Violation Rate [%]
1 agent + 2 obst.	PARM	1	746	100	8.5	22.1	18.5	9.4	0
		6	1388	100	7.3	25.5	15.1	10.9	0
	PARM*	1	1462	100	7.0	26.4	31.1	0	0
		6	3636	100	9.1	20.6	26.0	0	0
	PRIMER	6	140	100	4.2	22.0	33.0	0	0
3 agents + 2 obst.	PARM	1	728	100	25.7	17.1	36.0	3.2	0
		6	1415	100	9.3	21.0	49.0	7.8	0
	PARM*	1	1839	60	8.5	24.0	63.3	0	0
		6	5762	90	14.1	24.3	101.1	0	0
	PRIMER	6	109	100	6.6	22.6	60.2	0.3	0

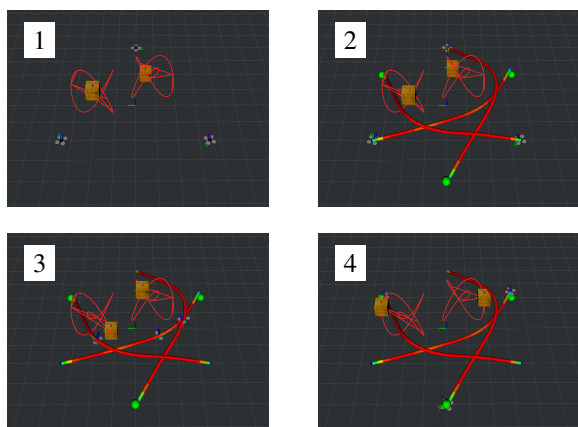


Fig. 4. Student multi-agent, multi-obstacle, simulation result: We made three imitation learning-based (student) agents fly around two dynamic obstacles. They started at the top-right corner and was commanded to fly to the down-left. For simplicity, we omitted FOV tripods visualization.

which were perception-aware, decentralized, asynchronous multiagent trajectory planners that enabled teams of agents to navigate uncertain environments while avoiding obstacles and deconflicting trajectories using perception information. Although these methods achieved state-of-the-art performance, they suffered from high computational costs, making it difficult for agents to replan at high rates.

To overcome this challenge, we presented PRIMER, a learning-based planner that was trained with imitation learning (IL) using PARM\* as the expert demonstrator. PRIMER is a computationally-efficient deep neural network and achieved a computation speedup of up to 5614 times faster than optimization-based approaches, while maintaining high performance. This speedup enables scalability to a large number of agents, making PRIMER a promising approach for large-scale swarm coordination.

Moving forward, our future work will focus on larger-scale simulations and hardware flight experiments to demonstrate the scalability and performance of PRIMER in complex environments with many agents and obstacles.

## 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, J. Tordesillas, R. Figueroa, J. Rached, J. Merkel, P. C. Lusk, and J. P. How, “Robust mader: Decentralized and asynchronous multiagent trajectory planner robust to communication delay,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 1687–1693.
- [6] 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,” *IEEE Robotics and Automation Letters*, 2023.
- [7] 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.
- [8] 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.
- [9] D. R. Robinson, R. T. Mar, K. Estabridis, and G. Hewer, “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.
- [10] 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.
- [11] J. Hou, X. Zhou, Z. Gan, and F. Gao, “Enhanced decentralized autonomous aerial swarm with group planning,” *ArXiv*, vol. abs/2203.01069, 2022.
- [12] 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.
- [13] C. Toumich, “Decentralized multi-agent planning for multirotors: a fully online and communication latency robust approach,” *arXiv preprint arXiv:2304.09462*, 2023.
- [14] 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.
- [15] 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.

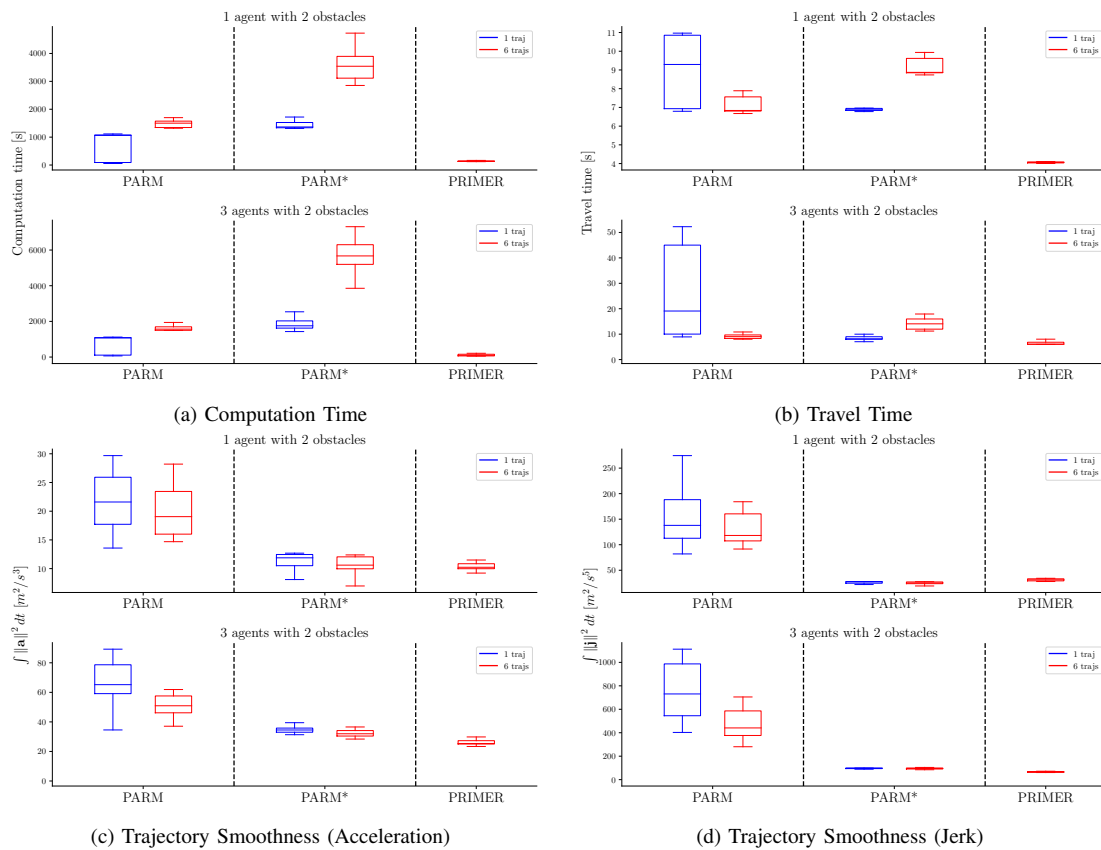


Fig. 5. Results of flight simulations. (a) The student’s computation time is much faster than that of the expert, and (b) the student’s travel time is also much shorter; this is mainly because of the faster computation time. (c-d) Since the student achieves faster replanning, it does not need to stop as the expert does, and that leads to smoother trajectory generation.

[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] C. Toumieh and D. Floreano, “High-speed motion planning for aerial swarms in unknown and cluttered environments,” *arXiv preprint arXiv:2402.19033*, 2024.

[18] 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.

[19] 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.

[20] K. Kondo, C. T. Tewari, M. B. Peterson, A. Thomas, J. Kinari, A. Tagliabue, and J. P. How, “Puma: Fully decentralized uncertainty-aware multiagent trajectory planner with real-time image segmentation-based frame alignment,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 13 961–13 967.

[21] 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.

[22] R. Van Parys and G. Pipeleers, “Distributed model predictive formation control with inter-vehicle collision avoidance,” in *ASCC*. IEEE, 2017.

[23] C. E. Luis, M. Vukosavljev, and A. P. Schoellig, “Online trajectory generation with distributed model predictive control for multi-robot motion planning,” *IEEE RA-L*, vol. 5, no. 2, pp. 604–611, 2020.

[24] C. Toumieh and A. Lambert, “Decentralized Multi-Agent Planning Using Model Predictive Control and Time-Aware Safe Corridors,” *IEEE RA-L*, vol. 7, no. 4, pp. 11 110–11 117, Oct. 2022.

[25] 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.

[26] I. Spasojevic, V. Murali, and S. Karaman, “Perception-aware time optimal path parameterization for quadrotors,” in *2020 ICRA*, 2020, pp. 3213–3219.

[27] J. Tordesillas and J. P. How, “Panther: Perception-aware trajectory planner in dynamic environments,” *IEEE Access*, vol. 10, pp. 22 662–22 677, 2022.

[28] 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.

[29] 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)*. IEEE, 2018, pp. 1–8.

[30] K. Kondo, A. Tagliabue, X. Cai, C. Tewari, O. Garcia, M. Espitia-Alvarez, and J. P. How, “Cgd: Constraint-guided diffusion policies for uav trajectory planning,” *arXiv preprint arXiv:2405.01758*, 2024.

[31] 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.

[32] A. Tagliabue, D.-K. Kim, M. Everett, and J. P. How, “Demonstration-efficient guided policy search via imitation of robust tube mpc,” *ICRA*, pp. 462–468, 2021.

[33] B. Park and H. Oh, “Vision-based obstacle avoidance for uavs via imitation learning with sequential neural networks,” *International Journal of Aeronautical and Space Sciences*, vol. 21, pp. 768 – 779, 2020.

[34] D. Mellinger and V. Kumar, “Minimum snap trajectory generation and control for quadrotors,” in *2011*, 2011, pp. 2520–2525.

[35] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[36] M. Everett, Y. F. Chen, and J. P. How, “Motion planning among dynamic, decision-making agents with deep reinforcement learning,” in *2018 IEEE/RSJ IROS*, 2018, pp. 3052–3059.

[37] 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.
- [38] 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.
- [39] S. Ross, G. Gordon, and D. Bagnell, “A reduction of imitation learning and structured prediction to no-regret online learning,” in *Proceedings of the fourteenth international conference on artificial intelligence and statistics. JMLR Workshop and Conference Proceedings*, 2011, pp. 627–635.
- [40] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.