

# Nearest-Neighbor-based Collision Avoidance for Quadrotors via Reinforcement Learning

Ramzi Ourari\*, Kai Cui\*, Ahmed Elshamanhory and Heinz Koepl

**Abstract**—Collision avoidance algorithms are of central interest to many drone applications. In particular, decentralized approaches may be the key to enabling robust drone swarm solutions in cases where centralized communication becomes computationally prohibitive. In this work, we draw biological inspiration from flocks of starlings (*Sturnus vulgaris*) and apply the insight to end-to-end learned decentralized collision avoidance. More specifically, we propose a new, scalable observation model following a biomimetic nearest-neighbor information constraint that leads to fast learning and good collision avoidance behavior. By proposing a general reinforcement learning approach, we obtain an end-to-end learning-based approach to integrating collision avoidance with arbitrary tasks such as package collection and formation change. To validate the generality of this approach, we successfully apply our methodology through motion models of medium complexity, modeling momentum and nonetheless allowing direct application to real world quadrotors in conjunction with a standard PID controller. In contrast to prior works, we find that in our sufficiently rich motion model, nearest-neighbor information is indeed enough to learn effective collision avoidance behavior. Our learned policies are tested in simulation and subsequently transferred to real-world drones to validate their real-world applicability.

## I. INTRODUCTION

The global market for unmanned aerial vehicles (UAVs) remains in the development stage with a potential market of \$45 billion in the civil sector alone [1], showing the enormous economic potential of drone swarms. Teams of robots and in particular quadcopters are found to be useful in many applications such as search-and-rescue missions [2], emergency communication [3] or package delivery [4], the reason for which lies at least partially in their potential for low-level distributed processing applications such as decentralized object tracking and detection. At the same time, the usage of drone swarms in the real world adds various additional coordination challenges [5]. As UAVs in swarms will often operate in close vicinity of each other, an important challenge for successful autonomous application of drone swarms is decentralized collision avoidance.

The study of collision avoidance has a long history and begins with traditional path-planning and sensor-based methods. Existing collision avoidance algorithms based on path-planning typically require expensive planning as well as full state information, complicating their usage in large

\*Authors contributed equally.

This work has been co-funded by the LOEWE initiative (Hesse, Germany) within the emergenCITY center and the Distr@l-project BlueSwarm (Project 71585164) of the Hessian Ministry of Digital Strategy and Development. The authors are with the Department of Electrical Engineering, Technische Universität Darmstadt, 64287 Darmstadt, Germany. Contact: {kai.cui, ahmed.elshamanhory, heinz.koepl}@bcs.tu-darmstadt.de

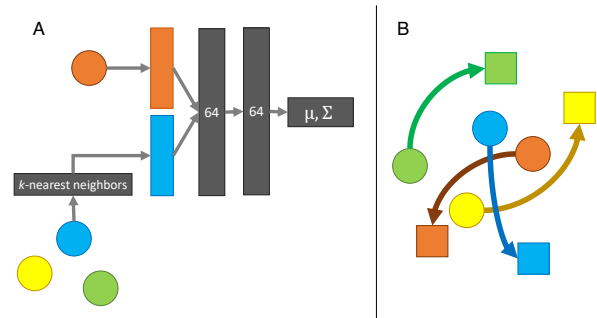


Fig. 1. Illustration of our proposed architecture (BICARL). A: The agent’s own observations (orange block) are concatenated with observations from the  $k$  nearest neighbors (blue block,  $k = 1$ ) and passed to a feed-forward neural network to obtain an immediate action distribution. B: The corresponding situation where agents need to negotiate their way toward their target while avoiding other dynamical agents. Each agent conditions their actions on observations of its  $k$ -nearest neighbors and their own state.

drone swarms [6]. On the other hand, although decentralized sensor-based methods can achieve very good performance – e.g. ORCA [7] – the disadvantage of such methods lies in their rigidity. While such algorithms may provide provable safety for the assumed model, they often lead to deadlock situations. In this context, end-to-end reinforcement learning (RL) may be able to provide better performance both in terms of deadlock occurrence and speed [8].

In this work, we propose a scalable, biologically well-motivated design to learning collision avoidance using end-to-end RL, i.e. Biologically-Inspired Collision Avoidance using RL (BICARL). In our method, we apply the biological insight into flocks of starlings (*Sturnus vulgaris*) interacting only with a limited number of their nearest neighbors to quadcopters in order to achieve decentralized collision avoidance that avoids crashes or getting stuck while learning successful policies in 30 minutes on two commodity desktop computers using an Intel Core i7-8700K CPU, 16 GB of RAM and no GPUs. This provides a scalable end-to-end learning approach to decentralized collision avoidance, reaching the performance of other state-of-the-art algorithms. Even though our motion model is more complex than models in prior works [9], [10], we find that the added complexity allows improved collision avoidance performance and at the same time direct application to real world quadrotors by combining with conventional low-level controllers. As a result, we obtain a very practical deep RL approach to decentralized collision avoidance that remains scalable and applicable to arbitrary task specifications while requiring

minimal information about the swarm, i.e. only from the nearest neighbor. Finally, we validate our algorithm both in simulation and real world application, verifying real world applicability.

## II. RELATED WORK

Traditional methods for collision avoidance include planning approaches with full state information [11], or hand-engineered sensor-based approaches using velocity obstacles [12] and potential fields [13], see also [14] for a survey.

Though the field of planning remains an active area of research [15], typical disadvantages of planning-based approaches are the requirement of full state information and high computational cost, barring scalable application in decentralized drone swarms. On the other hand, sensor-based methods may be applied in decentralized swarms, but have the drawback of potential deadlocks. Amongst the sensor-based methods, Optimal reciprocal collision avoidance (ORCA) is a method based on velocity obstacles [7], while force-based motion planning (FMP) is a state-of-the-art method based on potential fields [16].

*a) Learning-based approaches:* Recently, there have been many attempts to use machine learning for collision avoidance [8], [17], typically outperforming traditional state-of-the-art methods in terms of success rate and speed by learning a collision avoidance rule via reinforcement learning [18], though many prior works focus on simplified dynamics models and consider application only to small swarms or single drones [19].

The GA3C-CADRL algorithm [10] applies LSTMs in a single integrator model to learn a reaction to the position and velocities of all other agents in order to jointly avoid collisions. From the imitation learning realm, GLAS [20] focuses on imitating decisions of a global privileged planner with safety module. [21] extend GA3C-CADRL to 3D dynamics and fuses FMP [16] with RL in a hybrid algorithm. Finally, [22] applies imitation learning on ORCA for initialization, and then refines using RL. All of the aforementioned approaches except for [9] remain conditioned on information of all agent's (relative) positions and velocities. Similar to our work, [9] proposes RL-based collision avoidance via nearest neighbor information in the single integrator model, though their approach remains limited to very small swarms and in [10] was found to become stuck in large swarms.

*b) Biological inspiration:* In the field of robot navigation, there exist a great number of works inspired by biology. To name a few, one can find navigation algorithms inspired by bugs [23], optical flow navigation inspired by honey bees [24] or rule-based swarming models inspired by collective motion [25], see also [14] for a review.

In our work, we give a biological motivation for the nearest neighbor method. To be precise, we take inspiration from the behavior of flocks of starlings following a topological nearest-neighbor interaction rule in order to achieve robust swarm flight behavior [26]. In prior works, this type of biological insight has inspired flocking control design [27] and real-world deployment of drone swarms [28]. Somewhat

more related, [29] implements end-to-end-learned flocking control based on classical swarming models. However, their focus remains on flocking, and their observation model is fully observed and therefore not decentralized.

## III. MODEL

In this work we consider partially observable stochastic games (POSG) as the typical setting of multi-agent reinforcement learning. Formally, in our setting a POSG is a tuple  $(I, X, U, Z, T, r, p, \gamma)$ . The index set  $I = \{1, \dots, N\}$  is the set of agents,  $X$  is the state space,  $U = \bigotimes_{i \in I} U_i$  is the joint action space and  $Z = \bigotimes_{i \in I} Z_i$  is the joint observation space, where  $U_i$  and  $Z_i$  denote the action and observation spaces of each agent respectively. In our work, we use a deterministic transition function  $T: X \times U \rightarrow X$  and a random observation emission density  $p: Z \times X \rightarrow \mathbb{R}_{\geq 0}$ . The reward function  $r: X \times U \rightarrow \mathbb{R}^N$  and the discount factor  $\gamma \in (0, 1)$  give rise to the maximization objective

$$J_i(\pi) = \mathbb{E}_\pi \left[ \sum_{t=0}^{\infty} \gamma^t r_i(x_t, u_t) \mid x_0 \sim \mu_0 \right] \quad (1)$$

of agent  $i$  with initial state distribution  $\mu_0$  over joint policies  $\pi = (\pi^1, \dots, \pi^N)$  with  $\pi^i: U \times Z_i \rightarrow \mathbb{R}_{\geq 0}$  and

$$z_t \sim p(\cdot \mid x_t), \quad u_t^i \sim \pi^i(\cdot \mid z_t^i), \quad x_{t+1} = T(x_t, u_t) \quad (2)$$

for  $t \geq 0, i \in I$  and  $u_t \equiv (u_t^1, \dots, u_t^N), z_t \equiv (z_t^1, \dots, z_t^N)$ .

### A. Dynamics

In our work, we will consider both a 2D and a 3D case. The perhaps simplest model studied in multi-agent collision avoidance is the 2D single integrator model used in most prior works such as [7], [10], where the states  $x_t \equiv (\mathbf{p}_t^i, \mathbf{p}_t^{i,*})_{i \in I}$  consist of the  $\mathbb{R}^2$ -valued position  $\mathbf{p}_t^i = (x_t^i, y_t^i)$  as well as goal position  $\mathbf{p}_t^{i,*} = (x_t^{i,*}, y_t^{i,*})$  of agents  $i$ . As actions  $u_t^i \equiv \mathbf{v}_t^i$ , each agent may choose their  $\mathbb{R}^2$ -valued velocity  $\mathbf{v}_t^i$  directly as an action under the constraint  $\|\mathbf{v}_t^i\|_2 \leq v_{\max}$ , leading to the deterministic transition function defined by

$$\mathbf{p}_{t+1}^i = \mathbf{p}_t^i + \Delta t \cdot \mathbf{v}_t^i \quad (3)$$

for time step size  $\Delta t \geq 0$ .

We propose a more complex model with some degree of momentum by using the following modified double integrator model, where the state  $x_t \equiv (\mathbf{p}_t^i, \mathbf{v}_t^i, \theta_t^i, \omega_t^i, \mathbf{p}_t^{i,*})_{i \in I}$  is given not only by the positions, but also by the  $\mathbb{R}^2$ -valued velocities  $\mathbf{v}_t^i$  as well as the  $\mathbb{R}$ -valued yaw angle and its associated angular rate  $\theta_t^i, \omega_t^i$ .

Note that this model can alternatively be understood as part of our algorithm on top of the single integrator model, i.e. the algorithm keeps track of any additional states and simulates the modified double integrator dynamics to choose velocities at every time step. Therefore, results for this more complex dynamics model are nonetheless applicable and comparable to the single integrator case.

An action  $u_t^i \equiv (\tilde{\mathbf{v}}_t^i, \tilde{\omega}_t^i)$  of agent  $i$  with chosen target velocity  $\tilde{\mathbf{v}}_t^i \in \mathbb{R}^2$  and target angular velocity  $\tilde{\omega}_t^i \in \mathbb{R}$  leads

to the deterministic update

$$\mathbf{p}_{t+1}^i = \mathbf{p}_t^i + \mathbf{v}_t^i \cdot \Delta t, \quad (4)$$

$$\mathbf{v}_{t+1}^i = \mathbf{v}_t^i - c_v (\mathbf{R}(\theta_t^i) \cdot \tilde{\mathbf{v}}_t^i - \mathbf{v}_t^i) \cdot \Delta t, \quad (5)$$

$$\theta_{t+1}^i = \theta_t^i + \omega_t^i \cdot \Delta t, \quad (6)$$

$$\omega_{t+1}^i = \omega_t^i - c_\omega (\tilde{\omega}_t^i - \omega_t^i) \cdot \Delta t \quad (7)$$

with  $\|\tilde{\mathbf{v}}_t^i\|_2 \leq v_{\max}$ ,  $|\tilde{\omega}_t^i| \leq \omega_{\max}$  and yaw rotation matrix

$$\mathbf{R}(\theta_t^i) = \begin{bmatrix} \cos \theta_t^i & -\sin \theta_t^i \\ \sin \theta_t^i & \cos \theta_t^i \end{bmatrix}. \quad (8)$$

Importantly, although the yaw angle is not strictly required, we add the yaw angle to empirically obtain significantly better performance, as the added model richness allows agents to implicitly save information without requiring e.g. recurrent policy architectures. This model of medium complexity will at the same time allow us to directly use the desired velocity output as a set point of more traditional low-level quadrotor controllers such as PID controllers.

Note that depending on the specific task, we add task-specific transitions for the goal position, see the section on experiments. For the 3D case, we simply add another  $z$ -coordinate to position and velocity that remains unaffected by the rotation matrix. It should be further noted that we observe no discernible difference when applying some small amount of noise to the dynamics.

### B. Observation model

We let the observations of an agent  $i$  be randomly given by the own position and velocity as well as the relative bearing, distance and velocity of other agents inside of the sensor range  $K > 0$  and the goal, i.e.

$$z_t^i \equiv (\hat{\mathbf{p}}_t^i, \hat{\mathbf{v}}_t^i, d_t^{i,*}, \phi_t^{i,*}, \{(d_t^{i,j}, \phi_t^{i,j}, \mathbf{v}_t^{i,j}) \mid j \neq i: \|\mathbf{p}_t^j - \mathbf{p}_t^i\|_2 \leq K\}) \quad (9)$$

where the observations are Gaussian distributed according to

$$\hat{\mathbf{p}}_t^i \sim \mathcal{N}(\mathbf{p}_t^i, \sigma_p^2), \quad \hat{\mathbf{v}}_t^i \sim \mathcal{N}(\mathbf{v}_t^i, \sigma_v^2), \quad (10)$$

$$d_t^{i,*} \sim \mathcal{N}(\|\mathbf{p}_t^{i,*} - \mathbf{p}_t^i\|_2, \sigma_d^2), \quad (11)$$

$$\phi_t^{i,*} \sim \mathcal{N}(\phi_t^{i,*}, \sigma_\phi^2), \quad (12)$$

$$d_t^{i,j} \sim \mathcal{N}(\|\mathbf{p}_t^j - \mathbf{p}_t^i\|_2, \sigma_d^2), \quad (13)$$

$$\phi_t^{i,j} \sim \mathcal{N}(\phi_t^{i,j}, \sigma_\phi^2), \quad (14)$$

$$\mathbf{v}_t^{i,j} \sim \mathcal{N}(\mathbf{v}_t^j - \mathbf{v}_t^i, \sigma_v^2) \quad (15)$$

with noise standard deviations  $\sigma_p, \sigma_v, \sigma_d, \sigma_\phi > 0$  and bearing angles  $\phi_t^{i,*} = \arctan2(y_t^{i,*} - y_t^i, x_t^{i,*} - x_t^i)$ ,  $\phi_t^{i,j} = \arctan2(y_t^j - y_t^i, x_t^j - x_t^i)$  where  $\arctan2(y, x)$  is defined as the angle between the positive  $x$ -axis and the ray from 0 to  $(x, y)$ . Note further that the observations allow application of e.g. ORCA [7] and FMP [16]. In the 3D case, we additionally observe the  $z$ -coordinate difference to the target and other agents with Gaussian noise of standard deviation  $\sigma_p$ .

### C. Reward function

As the reward function for all of our experiments, for each agent  $i \in I$  we define a term for reaching the goal position and for avoiding collisions each by

$$r_i(x_t, u_t) = c_p \langle \mathbf{v}_t^i, \mathbf{p}_t^{i,*} - \mathbf{p}_t^i \rangle - c_c \sum_{j \in I \setminus \{i\}} \mathbf{1} \left( \|\mathbf{p}_t^j - \mathbf{p}_t^i\|_2 \leq C \right) \quad (16)$$

with desired target position  $\mathbf{p}_t^{i,*} \in \mathbb{R}^2$ , avoidance radius  $C \geq 0$  and reward / collision penalty coefficients  $c_p, c_c \geq 0$ , where  $\langle \cdot, \cdot \rangle$  denotes the standard dot product.

Note that this objective is selfish, i.e. agents are only interested in their own collisions and reaching their own goal. This manual choice of reward function alleviates the well-known multi-agent credit assignment problem in multi-agent RL [30], as shared, cooperative swarm reward functions are difficult to learn due to the increasing noise from other agent's behaviors as the number of agents rises.

## IV. METHODOLOGY

In this work, we propose a biologically-inspired design principle to learning collision avoidance algorithms. Recently, it was found that swarms of starlings observe only up to seven neighboring birds to obtain robust swarm flight behavior [26]. This implies that it should be sufficient to use a similar observation reduction for collision avoidance. Further, it is well-known that the multi-agent RL domain suffers from the combinatorial nature of multi-agent problems [30]. Hence, the reduction of observations to e.g. the closest  $k$  agents can greatly help learning effective behavior.

It is known that tractable exact solution methods and theoretical guarantees in the context of POSGs are sparse even in the fully cooperative case [31]. Instead, we apply independent learning, i.e. we ignore the non-stationarity of other agents and solve assumed separate single-agent problems via RL [32]. Furthermore, we share a policy between all agents via parameter sharing [33] and use the PPO algorithm to learn a single, shared policy [34]. For faster learning, we collect rollout trajectories from our simulation in parallel on two machines and use the RLlib [35] implementation of PPO. We implemented our environment in Unity ML-Agents [36].

We parametrize our policy by a simple feedforward network consisting of two hidden layers with 64 nodes, and ReLU activations except for the output layer which outputs parameters of a diagonal Gaussian distribution over actions. Actions are clipped after sampling to fulfill constraints. We preprocess observations by reducing to the information of the nearest  $k$  neighbors in  $L_2$  norm on the positions  $\mathbf{p}_t^i$ . All results in this work are for  $k = 1$ , i.e. we use observations  $(d_t^{i,j}, \phi_t^{i,j}, \mathbf{v}_t^{i,j})$  of the nearest neighbor  $j$ . Crucially, this observation reduction works well in our experiments, and it allows us to show that information of neighboring agents limited to the nearest neighbor information is indeed sufficient for collision avoidance. A resulting advantage of our design is that the resulting policy itself is very simple and memoryless.

TABLE I

HYPERPARAMETERS AND PARAMETERS USED IN ALL EXPERIMENTS.

Symbol	Function	Value
$\Delta t$	Time step size	0.02 s
$c_v$	Velocity coefficient	$1 \text{ s}^{-1}$
$c_w$	Angular rate coefficient	$1 \text{ s}^{-1}$
$\sigma_p, \sigma_v$	Noise standard deviations	1 mm, 10 mm
$\sigma_d, \sigma_\phi$	Noise standard deviations	1 mm, 0.1 mm
$C$	Avoidance radius	7 m
$c_p$	Reward coefficient	$0.3 \text{ s m}^{-2}$
$c_c$	Penalty coefficient	1
$v_{\max}$	Maximum velocity	$30 \text{ m s}^{-1}$
$\omega_{\max}$	Maximum angular velocity	$15 \text{ s}^{-1}$
$K$	Sensor range	10 m
$k$	Number of considered neighbors	1
$\gamma$	Discount factor	0.99
PPO		
$l_r$	Learning rate	$5 \times 10^{-5}$
$\beta$	KL coefficient	0.2
$\epsilon$	Clip parameter	0.3
$B$	Training batch size	50000
$B_m$	SGD Mini batch size	2500
$M$	SGD iterations	20

TABLE II

PERFORMANCE OF BICARL (OURS), ORCA [7], AND FMP [16] IN THE PACKAGE DELIVERY TASK, AVERAGED OVER 4 RUNS OF 100 SECONDS.

Test setup	Average collected packages		
# agents	BICARL	ORCA	FMP
4	<b>15.25</b>	1.5	3
6	<b>8.66</b>	1.16	2
8	<b>7.62</b>	1.25	2.12
10	<b>5.8</b>	0	1.2
12	<b>4.41</b>	0	0.41
14	<b>4.14</b>	0	0.25

Note that our policy architecture is very small and thus allows sufficiently fast computation on low-powered micro-controllers and UAVs such as the Crazyflie [37] used in our real world experiments. In comparison, previous works use Deep Sets [38] or LSTMs [10] to process a variable number of other agents, which will scale worse in large swarms and is more computationally costly to learn and apply.

The scenarios considered in this work include formation change and package delivery. In the formation change scenario, all drones start on a circle and have their goal set to the opposite side. In the package delivery scenario, drones are continuously assigned new random package locations and destinations after picking and dropping off packages, where the locations and destinations are uniformly distributed e.g. on a line, circle or in a rectangle.

## V. EXPERIMENTS AND RESULTS

For all of our experiments, we use the parameters and hyperparameters indicated in Table I. We apply curriculum learning [39] to improve learning performance, increasing the number of agents continuously while training. The single resulting policy is then evaluated for a varying number of

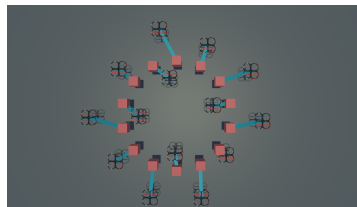


Fig. 2. A constellation where FMP is stuck in package delivery: At this number of agents, the computed forces nullify each other, resulting in freezing agents. The red squares visualize the current goal of agents and the blue lines signify which agent is assigned which goal.

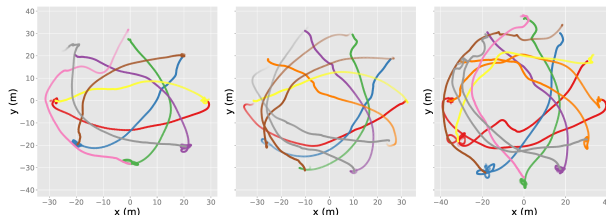


Fig. 3. Simulated trajectories of agents for 1-nearest neighbor BICARL with 2D dynamics. Left: 8 agents; Middle: 10 agents; Right: 12 agents.

agents to gauge generality and scalability. We consider a collision to occur when agents are closer than 1.5m and accordingly tuned the penalty radius during training to  $C = 7$  m to achieve reasonable behavior for up to 50 agents.

### A. Simulation Results

We have trained on two commodity desktop computers equipped with an Intel Core i7-8700K CPU, 16 GB RAM and no GPUs. We successfully and consistently learned good policies in 30 minutes upon reaching e.g.  $1.4 \cdot 10^6$  time steps. The reason for the fast training lies in the parallelization of collecting experience in Unity, as we can run multiple simulations at the same time to collect data, and the simulation in Unity is fast. We first compare our results to the algorithms FMP [16] and ORCA [7]. As mentioned earlier, for comparison we may also consider our dynamics as part of the algorithm for simpler single integrator dynamics.

a) *Learned collision avoidance behavior:* In the package delivery scenario, we place the agents on a circle of fixed radius and gradually increase the number of agents. As soon as agents become closer than 3.5m to the goal, we sample a new goal on the circle. In Table II, it can be observed that the average packages collected per drone decrease as the number of drones increases, as the drones will increasingly be in each other's way. Further, for the package delivery task, FMP and ORCA eventually run into a deadlock. An example for FMP can be seen in Fig. 2. In this case, our methodology provides robust, collision-free behavior that avoids getting stuck.

In the formation change scenario, during training we start with 4 diametrically opposed agents on a circle of radius 70m and gradually increase the size of the swarm up to 40. It can be seen in Fig. 3 that rotation around a fixed direction emerges. Increasing the number of agents leads to situations where ORCA and FMP get stuck. [40] demonstrates that the solution is cooperative collision avoidance. In line with

TABLE III

PERFORMANCE ANALYSIS OF BICARL (OURS), ORCA [7], AND FMP [16] IN CIRCLE FORMATION CHANGE, AVERAGED OVER 5 RUNS.

Test setup	Extra time to goal (s)			Minimal distance to other agents (m)			Extra travelled distance (m)		
# agents	BICARL	ORCA	FMP	BICARL	ORCA	FMP	BICARL	ORCA	FMP
5	4.02	14.52	<b>1.86</b>	<b>12.38</b>	4.94	6.18	1.28	1.17	<b>1.11</b>
10	4.51	15.31	<b>1.95</b>	<b>7.01</b>	4.78	3.78	1.26	1.65	<b>1.15</b>
15	4.87	14.98	<b>2.15</b>	<b>6.14</b>	3.42	4.44	1.27	1.82	<b>1.19</b>
20	6.51	18.68	<b>2.34</b>	3.14	<b>4.74</b>	3.87	1.45	2.7	<b>1.26</b>
25	7.52	20.13	<b>2.28</b>	3.3	<b>4.7</b>	3.34	1.60	3.56	<b>1.23</b>
30	7.42	31.51	<b>3.66</b>	3.52	<b>4.7</b>	2.94	1.61	4.68	<b>1.38</b>
35	<b>7.81</b>	41.15	N/A	2.52	<b>4.72</b>	N/A	<b>1.60</b>	5.71	N/A
40	<b>9.18</b>	45.64	N/A	2.35	<b>2.75</b>	N/A	<b>1.79</b>	10.17	N/A
45	<b>8.91</b>	76.25	N/A	1.94	<b>3.21</b>	N/A	<b>1.76</b>	8.14	N/A
50	<b>10.1</b>	81.03	N/A	1.67	<b>2.66</b>	N/A	<b>1.87</b>	17.95	N/A

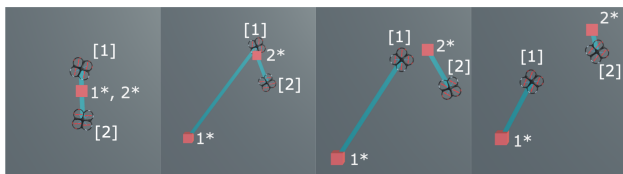


Fig. 4. Simulated flight behavior of two agents in the package delivery scenario with overlapping goals ( $1^*$ ,  $2^*$ ). In our approach, one drone yields and the other drone collects, obtaining a new goal. In FMP, the drones become stuck, while in BICARL one drone yields for the other drone. The red squares and blue lines have the same meaning as in Fig. 2. Time progresses from left to right.

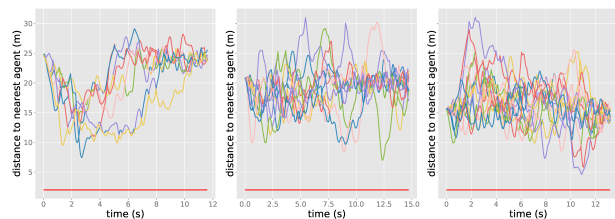


Fig. 5. Simulated minimum inter-agent distances achieved in the circle formation change scenario. The red line indicates the radius considered as collision. Left: 10 agents; Middle: 15 agents; Right: 20 agents.

this finding, our learned policy is able to capture such cooperation, i.e. one drone gives way to another as can be seen in Fig. 4 and supplementary videos.

Overall, from Table III we see that our solution is competitive, especially for many agents. Note that extra time to goal and travelled distance are measured relative to the baseline where drones fly in straight lines to the goal. Although FMP achieves very good results for small numbers of agents, at more than 35 agents FMP becomes stuck, while our method learns a robust behavior that works for large numbers of agents. Here, we tuned FMP parameters to obtain reasonably smooth flight. While improved results in FMP could be possible, additional parameter tuning would be required.

In Fig. 5, we find that our method successfully avoids collisions with other agents while reaching the goals. During training, collisions may be caused by nearby other agents

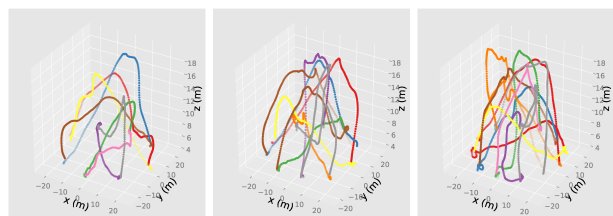


Fig. 6. Simulated trajectories of agents for 1-nearest neighbor BICARL with 3D dynamics. Agents have free space above their starting position and use it to navigate past each other. Left: 8 agents; Middle: 10 agents; Right: 12 agents.

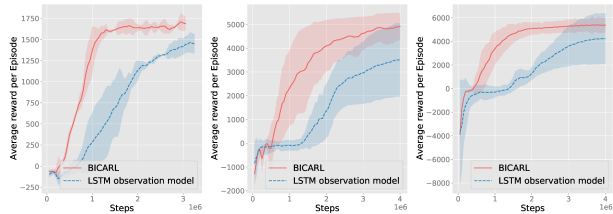


Fig. 7. Learning curve: Average achieved sum of rewards per episode over total time steps, plus variance as shaded region for three random seeds. It can be seen that our model learns significantly faster. Left: 2 agents; Middle: 4 agents; Right: 6 agents.

regardless of the own behavior. As a result, this often provides a negative feedback signal even if the drone itself is not responsible for the collision, resulting in behavior where agents avoid other agents even when the punishment for violating the avoidance radius (here  $C = 7$ ) is far off.

Analogously to the 2D case, we can consider formation change in the 3D case. In Fig. 6 the trajectories of our learned policy are depicted. It can be seen that our methodology remains successful in guiding the agents to their goal destination. Furthermore, to make use of the additional space above, agents begin flying above and below each other.

*b) Comparison to LSTM-based policies:* We compare our observation model with the LSTM variant proposed in [10]. We use an LSTM observation filter in conjunction with the two-layer architecture and run several tests with different

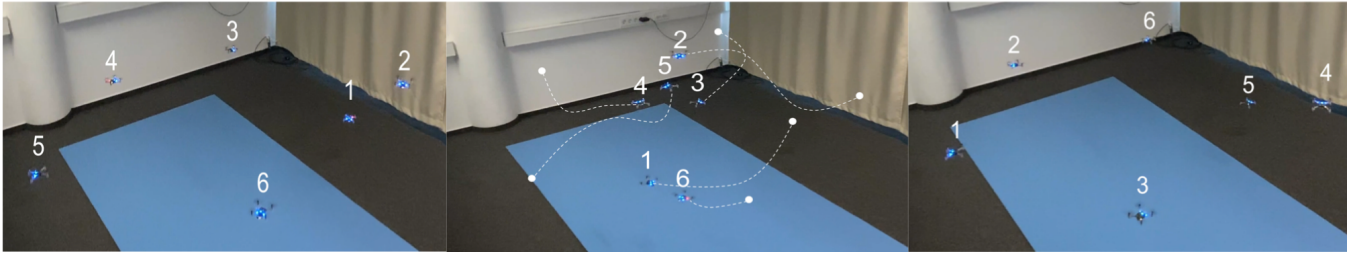


Fig. 8. Real world circle formation exchange with the policy for the 3D case embedded on a swarm of 6 Crazyflie nano-quadcopters. The drones successfully switch their positions to the antipodal points on a circle by engaging in 3D collision avoidance. Time progresses from left to right.

parameters (e.g. number of hidden units) for 2, 4 and 6 agents, comparing the average achieved reward in each case.

The learning curves in Fig. 7 suggest that our observation model achieves comparable results and faster learning. As expected, training LSTMs is slow and our method speeds up learning greatly, especially for many agents. While this does not exclude that in general and for improved hyperparameter configuration, the LSTM observation model will be superior, our method regardless has less parameters to fine-tune, has a constant computation cost regardless of the number of agents and hence is cheaper to both train and apply in the real world.

Although in [10] it has been noted that learned collision avoidance algorithms such as [9] based on neighborhood information tend to get stuck in large swarms or cause collisions, we obtain a seemingly opposed result. The reason is that our double integrator model together with the added yaw angle allows agents to avoid getting stuck by implicitly saving information or communicating with each other via the yaw angle. Since the problem at hand is multi-agent, a general optimal solution should act depending on the whole history of observations and actions. If two drones are stuck, they can therefore remember past information and change their reaction, leading to better results. Indeed, in our experiments we found that the yaw angle is crucial to obtaining good performance, leading us to the conclusion that a sufficiently rich motion model may in fact ease the learning of good collision avoidance behavior.

### B. Real World Experiments

We validate the real-world applicability of our policies by performing experiments on swarms of micro UAVs. We downscale lengths and speeds by 10 (since the used indoor drones are only 0.1m in length). We directly apply the desired velocity as a set point for the low-level controller instead of the acceleration as in conventional double integrator models, which is not easily regulated due to the non-linearity of typical quadrotor dynamics.

*a) Hardware Setup:* For our real-world experiments, we use a fleet of Crazyflie nano-quadcopters [37] with indoor positioning via the Lighthouse infrared positioning system and extended Kalman filter [41] with covariance correction [42]. Since the Crazyflies cannot sense each other, we simulate the inter-drone observations by exchanging information via TDMA-based peer-to-peer radio communication to broadcast the estimated position and velocity of each drone.

Although the sensor range  $K$  in the simulation can model a limited range of wireless communication, in our real indoor experiments there is no such limitation. Instead, drones ignore all information other than their supposed observations. As low-level controller, we use the default cascaded PID controller of the Crazyflie, setting desired velocities and yaw rates from the output of our policy directly as set point of the controller. Note that our code runs completely on-board, except for starting and stopping experiments.

*b) Evaluation:* We find that the policies obtained in simulation work as expected for two to six drones in reality, and expect similar results for more real drones. In Fig. 8, it can be seen that the drones successfully complete formation changes. Using our modified double integrator model, we easily achieve stable results in real world tests, as the cascaded PID low-level-controller is able to achieve desired velocities while stabilizing the system. Due to the simulation-to-reality gap stemming from model simplifications, we can observe slight control artefacts (oscillations). Nonetheless, the combination of traditional controller with learned policy successfully realizes the desired collision avoidance behavior, and the on-board inference time of the policy remains below 1 ms, with 30 $\mu$ s on our PCs. As a result, by using our motion model of medium complexity, we obtain a practical approach to decentralized quadrotor collision avoidance.

## VI. CONCLUSION

To summarize, in this work we have demonstrated that information of  $k$ -nearest neighbors, non-recurrent policies and a sufficiently rich motion model are enough to find robust collision avoidance behavior. We have designed a high-level model that allows application of the learned policy directly to the real world in conjunction with e.g. a standard cascaded PID controller. In the large swarm case and for collecting packages, our scalable and decentralized methodology appears to show competitive performance. Furthermore, our method avoids getting stuck and is very fast to learn.

Interesting future work could be to consider external static or dynamic obstacles such as walls. One could combine our approach for  $k > 1$  with Deep Sets [38], mean embeddings [43] or attention mechanisms [44] to further improve learning behavior. Finally, it may be of interest to further reduce the simulation-to-reality gap by using a more detailed model that models e.g. a 12-state quadrotor dynamics with simulated PID controller, motor lag, detailed drag models etc.

## REFERENCES

- [1] I. Kovalev, A. Voroshilova, and M. Karaseva, "Analysis of the current situation and development trend of the international cargo uavs market," in *Journal of Physics: Conference Series*, vol. 1399, no. 5. IOP Publishing, 2019, p. 055095.
- [2] Y. Karaca, M. Cicek, O. Tatli, A. Sahin, S. Pasli, M. F. Beser, and S. Turedi, "The potential use of unmanned aircraft systems (drones) in mountain search and rescue operations," *The American journal of emergency medicine*, vol. 36, no. 4, pp. 583–588, 2018.
- [3] D. Câmara, "Cavalry to the rescue: Drones fleet to help rescuers operations over disasters scenarios," in *2014 IEEE Conference on Antenna Measurements & Applications (CAMA)*. IEEE, 2014, pp. 1–4.
- [4] H. Shakhatareh, A. H. Sawalmeh, A. Al-Fuqaha, Z. Dou, E. Almaita, I. Khalil, N. S. Othman, A. Khreishah, and M. Guizani, "Unmanned aerial vehicles (uavs): A survey on civil applications and key research challenges," *Ieee Access*, vol. 7, pp. 48 572–48 634, 2019.
- [5] G. Chmaj and H. Selvaraj, "Distributed processing applications for uav/drones: a survey," in *Progress in Systems Engineering*. Springer, 2015, pp. 449–454.
- [6] D. Mellinger, A. Kushleyev, and V. Kumar, "Mixed-integer quadratic program trajectory generation for heterogeneous quadrotor teams," in *2012 IEEE international conference on robotics and automation*. IEEE, 2012, pp. 477–483.
- [7] J. Alonso-Mora, A. Breitenmoser, M. Ruffi, P. Beardsley, and R. Siegwart, "Optimal reciprocal collision avoidance for multiple non-holonomic robots," in *Distributed autonomous robotic systems*. Springer, 2013, pp. 203–216.
- [8] P. Long, T. Fan, X. Liao, W. Liu, H. Zhang, and J. Pan, "Towards optimally decentralized multi-robot collision avoidance via deep reinforcement learning," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 6252–6259.
- [9] Y. F. Chen, M. Everett, M. Liu, and J. P. How, "Socially aware motion planning with deep reinforcement learning," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 1343–1350.
- [10] M. Everett, Y. F. Chen, and J. P. How, "Collision avoidance in pedestrian-rich environments with deep reinforcement learning," *IEEE Access*, vol. 9, pp. 10357–10377, 2021.
- [11] M. Hamer, L. Widmer, and R. D'andrea, "Fast generation of collision-free trajectories for robot swarms using gpu acceleration," *IEEE Access*, vol. 7, pp. 6679–6690, 2018.
- [12] P. Fiorini and Z. Shiller, "Motion planning in dynamic environments using velocity obstacles," *The International Journal of Robotics Research*, vol. 17, no. 7, pp. 760–772, 1998.
- [13] K. Sigurd and J. How, "Uav trajectory design using total field collision avoidance," in *AIAA Guidance, Navigation, and Control Conference and Exhibit*, 2003, p. 5728.
- [14] M. Hoy, A. S. Matveev, and A. V. Savkin, "Algorithms for collision-free navigation of mobile robots in complex cluttered environments: a survey," *Robotica*, vol. 33, no. 3, pp. 463–497, 2015.
- [15] W. Hönig, J. A. Preiss, T. S. Kumar, G. S. Sukhatme, and N. Ayanian, "Trajectory planning for quadrotor swarms," *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 856–869, 2018.
- [16] S. H. Semnani, A. de Ruiter, and H. Liu, "Force-based algorithm for motion planning of large agent teams," 2019.
- [17] D. Willemsen, M. Coppola, and G. C. de Croon, "Mambpo: Sample-efficient multi-robot reinforcement learning using learned world models," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 5635–5640.
- [18] Y. F. Chen, M. Liu, M. Everett, and J. P. How, "Decentralized non-communicating multiagent collision avoidance with deep reinforcement learning," in *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2017, pp. 285–292.
- [19] G. Kahn, A. Villafior, V. Pong, P. Abbeel, and S. Levine, "Uncertainty-aware reinforcement learning for collision avoidance," *arXiv preprint arXiv:1702.01182*, 2017.
- [20] B. Rivière, W. Hönig, Y. Yue, and S.-J. Chung, "Glas: Global-to-local safe autonomy synthesis for multi-robot motion planning with end-to-end learning," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4249–4256, 2020.
- [21] S. H. Semnani, H. Liu, M. Everett, A. de Ruiter, and J. P. How, "Multi-agent motion planning for dense and dynamic environments via deep reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3221–3226, 2020.
- [22] D. Wang, T. Fan, T. Han, and J. Pan, "A two-stage reinforcement learning approach for multi-uav collision avoidance under imperfect sensing," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3098–3105, 2020.
- [23] K. N. McGuire, G. de Croon, and K. Tuyls, "A comparative study of bug algorithms for robot navigation," *Robotics and Autonomous Systems*, vol. 121, p. 103261, 2019.
- [24] W. E. Green and P. Y. Oh, "Optic-flow-based collision avoidance," *IEEE Robotics & Automation Magazine*, vol. 15, no. 1, pp. 96–103, 2008.
- [25] T. Vicsek and A. Zafeiris, "Collective motion," *Physics reports*, vol. 517, no. 3-4, pp. 71–140, 2012.
- [26] G. F. Young, L. Scardovi, A. Cavagna, I. Giardina, and N. E. Leonard, "Starling flock networks manage uncertainty in consensus at low cost," *PLoS Comput Biol*, vol. 9, no. 1, p. e1002894, 2013.
- [27] C. Liu, M. Wang, Q. Zeng, and W. Huangfu, "Leader-following flocking for unmanned aerial vehicle swarm with distributed topology control," *Science China Information Sciences*, vol. 63, no. 4, pp. 1–14, 2020.
- [28] P. Petráček, V. Walter, T. Báča, and M. Saska, "Bio-inspired compact swarms of unmanned aerial vehicles without communication and external localization," *Bioinspiration & Biomimetics*, vol. 16, no. 2, p. 026009, 2020.
- [29] P. Zhu, W. Dai, W. Yao, J. Ma, Z. Zeng, and H. Lu, "Multi-robot flocking control based on deep reinforcement learning," *IEEE Access*, vol. 8, pp. 150397–150406, 2020.
- [30] P. Hernandez-Leal, B. Kartal, and M. E. Taylor, "A survey and critique of multiagent deep reinforcement learning," *Autonomous Agents and Multi-Agent Systems*, vol. 33, no. 6, pp. 750–797, 2019.
- [31] K. Zhang, Z. Yang, and T. Başar, "Multi-agent reinforcement learning: A selective overview of theories and algorithms," *Handbook of Reinforcement Learning and Control*, pp. 321–384, 2021.
- [32] M. Tan, "Multi-agent reinforcement learning: Independent vs. cooperative agents," in *Proceedings of the tenth international conference on machine learning*, 1993, pp. 330–337.
- [33] J. K. Gupta, M. Egorov, and M. Kochenderfer, "Cooperative multi-agent control using deep reinforcement learning," in *International Conference on Autonomous Agents and Multiagent Systems*. Springer, 2017, pp. 66–83.
- [34] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [35] E. Liang, R. Liaw, R. Nishihara, P. Moritz, R. Fox, K. Goldberg, J. Gonzalez, M. Jordan, and I. Stoica, "Rllib: Abstractions for distributed reinforcement learning," in *International Conference on Machine Learning*. PMLR, 2018, pp. 3053–3062.
- [36] A. Juliani, V.-P. Berges, E. Teng, A. Cohen, J. Harper, C. Elion, C. Goy, Y. Gao, H. Henry, M. Mattar et al., "Unity: A general platform for intelligent agents," *arXiv preprint arXiv:1809.02627*, 2018.
- [37] W. Giernacki, M. Skwierczyński, W. Witwicki, P. Wroński, and P. Kozierski, "Crazyflie 2.0 quadrotor as a platform for research and education in robotics and control engineering," in *2017 22nd International Conference on Methods and Models in Automation and Robotics (MMAR)*. IEEE, 2017, pp. 37–42.
- [38] M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Póczos, R. Salakhutdinov, and A. J. Smola, "Deep sets," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017, pp. 3394–3404.
- [39] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proceedings of the 26th annual international conference on machine learning*, 2009, pp. 41–48.
- [40] P. Trautman and A. Krause, "Unfreezing the robot: Navigation in dense, interacting crowds," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2010, pp. 797–803.
- [41] M. W. Mueller, M. Hamer, and R. D'Andrea, "Fusing ultra-wideband range measurements with accelerometers and rate gyroscopes for quadcopter state estimation," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 1730–1736.
- [42] M. W. Mueller, M. Hehn, and R. D'Andrea, "Covariance correction step for kalman filtering with an attitude," *Journal of Guidance, Control, and Dynamics*, vol. 40, no. 9, pp. 2301–2306, 2017.
- [43] M. Hüttenrauch, S. Adrian, and G. Neumann, "Deep reinforcement learning for swarm systems," *Journal of Machine Learning Research*, vol. 20, no. 54, pp. 1–31, 2019.

- [44] A. Manchin, E. Abbasnejad, and A. van den Hengel, "Reinforcement learning with attention that works: A self-supervised approach," in *International Conference on Neural Information Processing*. Springer, 2019, pp. 223–230.