



Distributed Reinforcement Learning for Flexible and Efficient UAV Swarm Control

Venturini, Federico; Mason, Federico; Pase, Francesco; Chiariotti, Federico; Testolin, Alberto; Zanella, Andrea; Zorzi, Michele

Published in:

IEEE Transactions on Cognitive Communications and Networking

DOI (link to publication from Publisher):

[10.1109/TCCN.2021.3063170](https://doi.org/10.1109/TCCN.2021.3063170)

Publication date:

2021

Document Version

Accepted author manuscript, peer reviewed version

[Link to publication from Aalborg University](#)

Citation for published version (APA):

Venturini, F., Mason, F., Pase, F., Chiariotti, F., Testolin, A., Zanella, A., & Zorzi, M. (2021). Distributed Reinforcement Learning for Flexible and Efficient UAV Swarm Control. *IEEE Transactions on Cognitive Communications and Networking*, 7(3), 955-969. Advance online publication. <https://doi.org/10.1109/TCCN.2021.3063170>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

Distributed Reinforcement Learning for Flexible and Efficient UAV Swarm Control

Federico Venturini*, Federico Mason*, Francesco Pase*,
Federico Chiariotti[‡], Alberto Testolin*[†], Andrea Zanella*, Michele Zorzi*

*Department of Information Engineering, University of Padova - Via Gradenigo, 6/b, Padova, Italy

[‡]Department of Electronic Systems, Aalborg University - Fredrik Bajers Vej 7C, Aalborg, Denmark

[†]Department of General Psychology, University of Padova - Via Venezia, 8, Padova, Italy

*†{venturini, masonfed, pasefrance, testolin, zanella, zorzi}@dei.unipd.it, [‡]fchi@es.aau.dk

Abstract—Over the past few years, the use of swarms of Unmanned Aerial Vehicles (UAVs) in monitoring and remote area surveillance applications has become widespread thanks to the price reduction and the increased capabilities of drones. The drones in the swarm need to cooperatively explore an unknown area, in order to identify and monitor interesting targets, while minimizing their movements. In this work, we propose a distributed Reinforcement Learning (RL) approach that scales to larger swarms without modifications. The proposed framework relies on the possibility for the UAVs to exchange some information through a communication channel, in order to achieve context-awareness and implicitly coordinate the swarm's actions. Our experiments show that the proposed method can yield effective strategies, which are robust to communication channel impairments, and that can easily deal with non-uniform distributions of targets and obstacles. Moreover, when agents are trained in a specific scenario, they can adapt to a new one with minimal additional training. We also show that our approach achieves better performance compared to a computationally intensive look-ahead heuristic.

Index Terms—Artificial intelligence, distributed decision making, mobile robots, neural network applications.

I. INTRODUCTION

The high data rate achievable with modern wireless communications and the increasing computational power of embedded systems, along with the sharp price reduction of commercial Unmanned Aerial Vehicles (UAVs), have enabled the use of swarms of drones for Smart City services [1]. Thanks to their size, flexibility and flight ability, these swarms represent a new solution for a plethora of different applications, such as remote surveillance, distributed sensing, wireless coverage extension and object tracking [2].

Over the past few years, researchers have proposed several UAV-based systems [3], but achieving an efficient distributed control is a complex problem, whose solution is often task-dependent. In this context, it is important to properly define the different sub-tasks of surveillance, monitoring, mapping and tracking [4]. In this work, we assume that targets are static, but occupy random positions in the monitored area. Moving UAVs are equipped with sensors that can detect targets within a limited sensing range, and a radio interface that makes it possible to share position information and sensing data. The UAVs need to coordinate to explore the area and find the targets without colliding with each other or with obstacles.

The problem of identifying fixed targets arises in several practical situations, ranging from the generation of real-time

flood maps [5] to the detailed tracking of weeds in agriculture [6], but an efficient initial exploration is of interest even for larger classes of problems, e.g., considering moving targets. One such example is wildfire monitoring in dry regions [7], which can be effective as long as the UAVs move faster than the spread of the fires.

The dynamic nature of these problems, in which actions can have long-term consequences and affect the future evolution of the environment in complex ways, makes them a natural application area for Reinforcement Learning (RL) techniques [8]. However, due to the curse of dimensionality, a centralized approach to the problem (i.e., using a single controller) is feasible only for very small swarms. In order to design a scalable system, Multi-Agent Reinforcement Learning (MARL) techniques need to be used, but the non-stationarity of the environment [9] complicates the system design and the agent training. This additional complexity makes MARL an open research field, and the different degrees of centralization and communication between agents make the configuration of the learning system an interesting problem to investigate.

In this work, we consider a MARL framework for exploration and surveillance. Our aim is to find a flexible Machine Learning (ML) strategy to explore and monitor a certain area with a swarm of UAVs that can exchange information within a certain coverage range. Performance is determined by the ability of the drones to find and reach the targets, which are located in unknown positions.

In our framework, the observations of other agents are shared through a radio channel and used to make decisions and to avoid collisions, thus encouraging cooperation. We define a Deep Q-Network (DQN) algorithm and demonstrate its efficiency with limited training, comparing it to a benchmark look-ahead heuristic and showing that our approach can better explore the environment and reach the targets faster. We also perform a transfer learning experiment, showing that agents trained on a certain map can learn to adapt to a completely new scenario much faster than restarting the training from scratch.

We adopted a general model, using a grid-world representation and making a limited number of assumptions on the nature of the task. Nonetheless, we show that our system can be implemented in several different scenarios. In particular, the map is not entirely visible to the UAVs, there are obstacles, and targets are in unknown positions (often clustered together, making clusters rarer and thus harder to find). These features

make MARL highly complex, especially when considering limited communication capabilities: to the best of our knowledge, our work is the first to apply it in such challenging conditions.

Our approach to solve the problem is to model the state as a series of correlated maps, which contain different information on the environment, making the learning framework extendable to even more complicated scenarios.

The contributions of this paper can be summarized as follows:

- We formulate a Networked Distributed Partially Observable Markov Decision Process (ND-POMDP) framework for swarm management in a complex environment and propose a MARL architecture to address such a problem;
- We show that the proposed system can outperform computationally heavy heuristics and transfer its knowledge to different scenarios with limited retraining;
- We analyze the effect of bigger changes in the environment, such as changing the size of the map or the number of drones, and show that transfer learning is still effective;
- We show that the system is robust to channel impairments, and can perform very well even in realistic scenarios that differ from the more abstract models used in the training phase.

A preliminary version of this paper was presented at ACM DroNet 2020 [10]; this version has a significantly updated system model, considering different map sizes and the presence of obstacles as well as a different MARL solution, and more extensive results on the performance of our approach. Moreover, we have added the analysis of the impact of the communication channel on the system's performance, and tested the proposed solution in a map obtained from real data.

The rest of the paper is divided as follows: first, Sec. II analyzes the related work in the field. The system model and MARL algorithm are presented in Sec. III. The experimental setup is reported in Sec. IV, while the experimental results, including transfer learning experiments, are reported in Sec. V. Finally, Sec. VI concludes the paper and presents some possible avenues of future work.

II. RELATED WORK

An extensive taxonomy of multi-agent solutions was presented in [11]. The general approaches adopted to solve the MARL problem can be cast into one of these four frameworks: (1) a single agent architecture that interacts with multiple copies of itself, generating emergent behaviors; (2) communication between agents of the same type with improved coordination; (3) cooperation between agents with different specialized goals achieving coordinated behavior; and (4) modeling other agents' behaviors and planning a response [12].

The authors in [13] study the first of these four approaches and use the tabular Q-learning algorithm to guide drones to survey an unknown area, showing that even the simplest MARL algorithm can improve the overall system rewards. Similarly, in [14] and [15] the MARL framework is applied to a more complex problem in which a UAV network is adopted

to provide flexible wireless communication. However, in these works the MARL algorithm is used to optimize resource allocation instead of guiding drones, so that a coordinated exploration strategy is missing.

An interesting research direction for MARL is pioneered in [16], which uses Deep Neural Networks (DNNs) to represent and learn more complex Q-functions [17]. At first, the authors study the performance of one network trained for all agents, which then share the same parameters during the execution phase (this is also our approach). A second proposed system uses the Differentiable Inter-Agent Learning (DIAL) framework, in which agents learn meaningful real-valued messages to be exchanged in order to improve cooperation: this allows for faster training, but the model is limited to a very small number of agents.

Other works use RL in the practical scenarios discussed above: in [5], the authors adopt a MARL approach to control a flood-finding swarm of UAVs. However, the model only considers a swarm with a fixed number of drones, and the experimental results are not compared to state-of-the-art heuristics. In [6], a reinforced random walk model is exploited to map weeds in an agricultural setting, taking noisy acquisition into account and solving the issue with collective observations. Random walks are then biased based on the positions of the already discovered targets, which have to be properly mapped, along with the distances from other drones in the network. In this case, the authors considered swarms of variable sizes, but the random walk needs to be manually tuned for each setting. Another recent study [7] considers wildfire spread monitoring, checking how the fire evolves and spreads in the map from a known starting point. The authors define the problem as a Decentralized Partially Observable Markov Decision Process (Dec-POMDP) [18] and carry out several experiments, as well as comparisons against a greedy heuristic (similar to the look-head method we studied in this work). A target-tracking application for disaster scenarios, with a model similar to our own but applied to a single drone, is described in [19]. Finally, [20] considers a MARL system with realistic communication, where a swarm of drones needs to get data from an Internet of Things (IoT) sensor network. This is a much simpler problem, as the position of the targets is known in advance, and the MARL framework only needs to optimize the trajectories.

The MARL approaches can also fit models in which UAV connectivity is important: in [21], a framework including RL and game theory is used to plan the path of two drones that need to save energy and minimize the interference to the ground network while maintaining a cellular connection. Furthermore, in [22] the authors design a centralized RL system to maximize coverage for a swarm of aerial base stations serving mobile users on the ground. A similar approach is taken in [23], which redefines the problem in terms of Quality of Experience (QoE) maximization for the users. For a fuller communication-oriented perspective on the use of RL for UAV networks, we refer the reader to [24].

These works have similar objectives to our own, but either go back to the single-agent setting or have restrictive assumptions: as an example, [7] considers well-known fire patterns,

Symbol	Description	Symbol	Description
\mathcal{M}	Coordinate set	\mathcal{O}	Observation space of the system ND-POMDP
M	Map grid size	Φ	Matrix of cell values
K	Number of targets	\mathbf{X}	Matrix of UAV positions
\mathbf{z}_k	Coordinates of the k -th target	Ω	Matrix of obstacle positions
σ	Standard dev. of the target Gaussian functions	$\hat{\Phi}$	Observed cell value matrix
$\phi(\cdot)$	Cell value function	$\hat{\Omega}$	Observed obstacle position matrix
\mathcal{U}	Set of UAVs	\mathbf{X}_u	Observed UAV position matrix for u
U	Number of UAVs (cardinality of \mathcal{U})	F	Observation window size (in number of cells)
d_{sparse}	Minimum target distance in the sparse scenario	ψ	Penalty for collisions
$\omega(\cdot)$	Obstacle location function	θ	Penalty for moving to forbidden areas
η	Fraction of the map occupied by obstacles	ρ	Obstacle value
ζ	Field of View of each UAV	$\nu_u(\mathbf{x}_u, \mathbf{a}_u)$	Invalid move indicator function for UAV u
h_{\min}	Minimum obstacle size (in number of cells)	$\chi_u(\mathbf{X}, \mathbf{A})$	Collision indicator function for UAV u
h_{\max}	Maximum obstacle size (in number of cells)	$r_u(s, \mathbf{a})$	Reward for UAV u
ℓ_i	Lower left corner coordinates of the i -th obstacle	π	Observation-action policy
\mathcal{H}_i	Set of cells occupied by the i -th obstacle	$R_{u,t}(\pi)$	Long-term reward for u using policy π
\mathbf{h}_i	Size of the i -th obstacle	γ	Exponential discount factor
N	Episode duration (steps)	e_t	Experience sample
\mathcal{S}	State space of the system ND-POMDP	α	Learning rate
$\mathcal{V}(s)$	Valid move space for state s	B_{size}	Size of a learning batch
\mathcal{A}	Action set	$Q(o_u, a_u)$	Q-value estimate of R
\mathbf{a}_u	Action for UAV u	n_q	Model update period steps)

TABLE I: Notation definitions.

which can be extensively learned, with a known starting point. In our case, the initial positions of the targets and of the UAVs are not the same across different episodes, making the model more general and complicating the learning task. Furthermore, unlike previous efforts in the literature, we exploit the transfer learning paradigm, showing how our model can easily adapt to scenarios with obstacles, realistic maps, and different swarm sizes. To the best of our knowledge, our work presents the most complex environment to date, in which a single architecture can deal with different map and swarm sizes, different numbers of targets to track, and the presence of obstacles.

III. SYSTEM MODEL

In the following, we first present the environment in which the UAVs operate. We give a full list of the notation used in Table I as a reference to the reader.

A. Environment

The system environment consists of a square grid of size $M \times M$. Each cell of the grid (we will refer to a cell or a location interchangeably in the following) is identified by its coordinates $\mathbf{m} \in \mathcal{M}$, where $\mathcal{M} = \mathcal{X} \times \mathcal{Y}$, and $\mathcal{X} = \mathcal{Y} = \{0, \dots, M-1\}$. We place a set of K targets on the map, which represent the objectives of the UAV surveillance application. The position of the k -th target is denoted as $\mathbf{z}_k = (x_k, y_k)$.

We then generate a set of K bivariate Gaussian functions over the grid, which represent the visibility of a target to the UAVs, with the same covariance matrix $\Sigma = \begin{pmatrix} \sigma^2 & 0 \\ 0 & \sigma^2 \end{pmatrix}$. The mean $\mathbf{z}_k = (z_{k,1}, z_{k,2})$ corresponds to the coordinates of the target. Note that the Gaussian functions do not represent actual distributions, but rather the full view of the UAVs, which can see a target from afar. The value of σ can be interpreted as the distance at which a target can be identified, as larger values of σ mean that the target is visible from further away.

Each cell can then be associated with a weight $\phi(\mathbf{m})$, which represents the *value* of the location, which increases with the

proximity to a target, and is given by the maximum of the Gaussian functions in that point, normalized in such a way that the target locations have values equal to 1:

$$\phi(\mathbf{m}) = \max_{k \in \{0, \dots, K-1\}} e^{-\frac{1}{2}((\mathbf{m}-\mathbf{z}_k)^T \Sigma (\mathbf{m}-\mathbf{z}_k))}. \quad (1)$$

If $\phi(\mathbf{m})$ is smaller than 0.01, it is set to 0, as the UAVs cannot see any target from that location. Under these conditions, the most valuable cells coincide with the center of each Gaussian function, which represents one of the targets in the considered scenario. While the environment is static, the UAVs move within the map with the aim of positioning themselves over the targets as fast as possible. We denote the set of UAVs by \mathcal{U} , and by U its cardinality.

In this work, we consider two different distributions for the targets, named *sparse* and *cluster*, which are characterized by different correlations among the target positions. In both cases, the first target is randomly placed on the grid following a 2D uniform distribution: \mathbf{z}_0 can take any value in \mathcal{M} with equal probability. The other targets are then placed sequentially, according to the following rules. In the sparse scenario, the position \mathbf{z}_i of the i -th target is randomly chosen in the set $\mathcal{M}_i^{\text{sparse}} = \{\mathbf{m} \in \mathcal{M} : \|\mathbf{m} - \mathbf{z}_j\|_2 > d_{\text{sparse}}, \forall j < i\}$, with probability mass distribution

$$P_{\text{sparse}}(\mathbf{z}_i = \mathbf{m}) = \frac{\|\mathbf{m} - \mathbf{z}_0\|_2}{\kappa_i^{\text{sparse}}}. \quad (2)$$

where $\kappa_i^{\text{sparse}} = \sum_{\mathbf{m} \in \mathcal{M}_i^{\text{sparse}}} \|\mathbf{m} - \mathbf{z}_0\|_2$ is a normalization factor. Hence, the other targets tend to be distributed far from the first, with a minimum distance d_{sparse} between each other.

In the cluster scenario, instead, the i -th target can take any position in the set $\mathcal{M}_i^{\text{cluster}} = \{\mathbf{m} \in \mathcal{M} : \|\mathbf{m} - \mathbf{z}_j\|_2 > 1, \forall j < i\}$ with probability mass distribution

$$P_{\text{cluster}}(\mathbf{z}_i = \mathbf{m}) = \frac{1}{(1 + \|\mathbf{m} - \mathbf{z}_0\|_2) \kappa_i^{\text{cluster}}}. \quad (3)$$

where $\kappa_i^{\text{cluster}} = \sum_{\mathbf{m} \in \mathcal{M}_i^{\text{cluster}}} \frac{1}{(1 + \|\mathbf{m} - \mathbf{z}_0\|_2)}$ is the normalization factor. In this case, the targets tend to cluster around the first

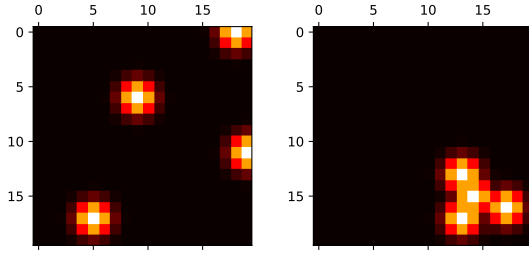


Fig. 1: Two examples of the sparse (left) and cluster (right) target distributions.

one, but cannot occupy adjacent cells, since the minimum distance must be greater than 1.

An example of the two target placements is shown in Fig. 1. These two distributions represent two plausible configurations of targets in tracking applications: in wildlife monitoring, some species of animals might tend to herd together, while more territorial ones will have a sparser distribution on the map. The same goes for a battlefield scenario, in which groups of soldiers might act together as a tight formation, while guerrilla-style fighting will involve a much sparser distribution of forces.

In a more complex version of the scenario, the map does not just have targets that the UAVs need to find and reach, but *obstacles* as well: in an urban scenario, these might be tall buildings or designated no-fly zones, while in a natural scenario they might correspond to natural obstacles such as boulders or tall trees. We define a function $\omega(\mathbf{m})$, which is equal to 1 if the cell corresponds to an obstacle and 0 otherwise. Then, we denote by η the portion of the map occupied by obstacles:

$$\eta = \sum_{\mathbf{m} \in \mathcal{M}} \frac{\omega(\mathbf{m})}{M^2}. \quad (4)$$

Cells inside an obstacle are considered impassable, like the map borders, and the UAVs that try to move on an obstacle will remain in the same cell.

For the training of our algorithm, we assumed that obstacles are rectangular and randomly scattered in the area. The i -th obstacle is determined by its dimensions \mathbf{h}_i and by the position of its lower left corner ℓ_i . We formally define the obstacle as the set \mathcal{H}_i :

$$\begin{aligned} \mathcal{H}_i &= \{\mathbf{m} = (m_1, m_2) \in \mathcal{M} : \\ &m_1 \in \{\ell_{i,1}, \dots, \ell_{i,1} + h_{i,1} - 1\}, \\ &m_2 \in \{\ell_{i,2}, \dots, \ell_{i,2} + h_{i,2} - 1\}\}. \end{aligned} \quad (5)$$

Obstacles are generated sequentially, like the targets, and for each obstacle i the dimensions \mathbf{h}_i are drawn uniformly from the set $\{h_{\min}, h_{\max}\} \times \{h_{\min}, h_{\max}\}$. The lower left corner position ℓ_i is then drawn from a uniform distribution in the set $\mathcal{M}_i^{\text{obs}}$, the subset of the map defined as:

$$\begin{aligned} \mathcal{M}_i^{\text{obs}} &= \{\ell \in \mathcal{M} : \mathcal{H}_i \subset \mathcal{M}, \|\mathbf{n}, \mathbf{z}_k\|_2 > 1, \\ &\forall \mathbf{n} \in \mathcal{H}_i, k \in \{0, \dots, K-1\}, \\ &d(\mathcal{H}_i, \mathcal{H}_j) \geq 2, \forall j < i\}, \end{aligned} \quad (6)$$

where $d(\mathcal{H}_i, \mathcal{H}_j) = \min_{\mathbf{m}_i \in \mathcal{H}_i, \mathbf{m}_j \in \mathcal{H}_j} \|\mathbf{m}_i - \mathbf{m}_j\|_2$ is the distance between the obstacles i and j . The three constraints force the obstacle to be entirely inside the map, not to be directly adjacent to any of the targets, and not to touch other obstacles. The choice of these constraints was motivated by the necessity to guarantee the existence of a clear path to the targets from any point in the map.

We consider multiple *episodes* of N steps: in each episode, the targets, UAVs, and obstacles are redistributed in the map, and the swarm must locate the targets in as few steps as possible. We consider discrete time slots, so that each drone can move by a single cell at each time step. Furthermore, we assume that a UAV has a limited Field of View (FoV), i.e., it can only know the value of the cells within a radius ζ . This framework allows us to represent many different applications and scenarios by changing the size of the grid, the number of drones, targets and obstacles, the FoV range ζ and the target visibility parameter σ . It can also be easily extended to dynamic targets.

At the beginning of each episode, each UAV only knows the values of the cells within the swarm's FoV. The drones assume that all unexplored points of the map are associated with the maximum $\phi(\mathbf{m})$. Then, each UAV moves independently at each time step n : as the swarm explores the environment, each drone discovers the values of the map locations that it has covered, and updates its information according to $\phi(\mathbf{m})$. We highlight that the knowledge about the map is instantly shared, which means that each drone receives the observations that all the other drones have acquired. This is always true during training, whereas in some testing episodes we also experiment the scenarios in which unreliable communications affect the shared messages. The objective of the swarm for each episode is to position each of its UAVs above a target as quickly as possible.

B. Communication model

We consider the swarm to only have partial observations: as the size of the map might be too large for the swarm to effectively coordinate over it, we consider each UAV to have up-to-date knowledge only inside the $F \times F$ square with it at the center, with $F \leq M$. If the distance between the UAV and the edge of the map is lower than F , the square will consider the edge of the map as the edge of the visible region, and the UAV will no longer be at its center, in order to avoid modeling the area outside the map. This assumption allows us to model communication constraints in the problem, as UAVs need to share the observed parts of the map with the other components of the swarm; however, F should not be confused with the FoV ζ , as the former represents the size of the portion of the map that each UAV considers when deciding its next action, while the latter represents the size of the portion of the map that the UAV can sense directly at each moment. In our case, we always have $F \geq \zeta$.

C. ND-POMDP formulation

The described scenario is modeled as an ND-POMDP [25], i.e., a Markov Decision Process (MDP) where the system state

in not directly observable and is influenced by the actions of multiple agents, whose behavior is not centrally coordinated. Indeed, the swarm only has limited knowledge of the map, and the UAVs can take actions independently and have independent rewards. We observe that ND-POMDP is a particular class of Decentralized POMDP (Dec-POMDP) for which not all agents interact with each other [26]. Convergence to the optimal solution for this kind of problem has been proven for classical reinforcement methods [27], although not for deep models: as most works in the literature, we will use a benchmark to evaluate the performance of our scheme. Formally, an ND-POMDP is identified by a 5-tuple, composed of a state space \mathcal{S} , an agent space \mathcal{U} , a joint action space \mathcal{A} , an observation space \mathcal{O} , and a reward map $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^U$, where $U = |\mathcal{U}|$.

The complete system state s is given by five matrices: a matrix for the current position of the UAVs, one matrix each for the map of the already discovered targets and obstacles, and one matrix each for the full map of targets and obstacles. The positions of the UAVs are contained in the $2 \times U$ matrix \mathbf{X} , while the features of the map are represented by the two $M \times M$ matrices Φ and Ω , which contain the value $\phi(\mathbf{m})$ of each cell and the function $\omega(\mathbf{m})$ representing the location of the obstacles. Clearly, the maps with the full view of targets and obstacles are not initially known by the UAVs, which will then need to explore the area.

Furthermore, the UAVs do not know the features of cells that have not been explored: the observed features of the map are contained in the $F \times F$ observed value matrix $\hat{\Phi}$, whose elements are equal to $\phi(\mathbf{m})$ if the cell has been explored and 1 otherwise, and the $F \times F$ observed obstacle matrix $\hat{\Omega}$, whose elements are equal to $\omega(\mathbf{m})$ if the cell has been explored and 0 otherwise. The observation $o_u \in \mathcal{O}$ that is available to drone u is then given by \mathbf{X}_u , $\hat{\Phi}_u$, and $\hat{\Omega}_u$, defined as the $F \times F$ subsets of \mathbf{X} , $\hat{\Phi}$ and $\hat{\Omega}$ centered in \mathbf{x}_u .

In our case, each UAV can either stay over the same cell or move to one of the four adjacent cells. However, obstacles are impassable in our environment definition, and the UAVs cannot move outside the map, so UAVs will simply stand in place if they attempt an action that violates the constraints. We define the action space $\mathcal{A} = \{(0, 0), (0, 1), (1, 0), (0, -1), (-1, 0)\}^U$. An action for the swarm is then a vector $\mathbf{a} \in \mathcal{A}$, which contains the individual UAVs' actions, denoted as \mathbf{a}_u for drone u . We first define function $\nu(\mathbf{x}_u, \mathbf{a}_u)$, which is 1 if the action is valid, i.e., it does not lead the UAV to fly outside the map or into an obstacle, and zero otherwise:

$$\nu(\mathbf{x}_u, \mathbf{a}_u) = \begin{cases} 1, & \text{if } \mathbf{x}_u + \mathbf{a}_u \in \mathcal{M} \wedge \omega(\mathbf{x}_u + \mathbf{a}_u) = 0; \\ 0, & \text{otherwise.} \end{cases} \quad (7)$$

The position of each drone is then updated in the following way:

$$\mathbf{x}_u(t+1) = \mathbf{x}_u(t) + \mathbf{a}_u(t)\nu(\mathbf{x}_u(t), \mathbf{a}_u(t)). \quad (8)$$

Fig. 2 shows an example of the system state at the beginning and in an advanced stage of an episode, with two drones and four targets located in a 20×20 map with no obstacles (in this case, we set $F = M = 20$). In particular, the drones' positions are shown on the left (in yellow), the observed value

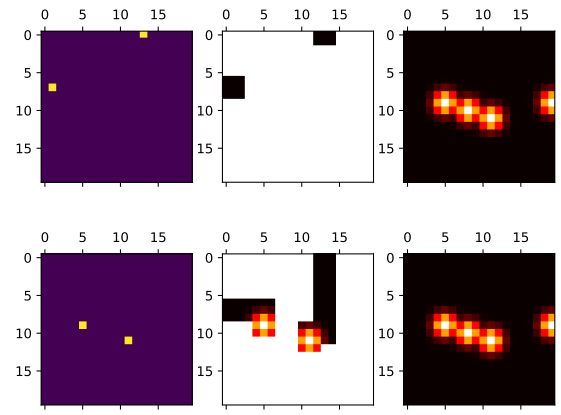


Fig. 2: Drone positions (left), known map (center), real map (right). Beginning (above) and end (below) of an episode.

map is in the center, and the real value map is on the right. In the figure, darker cells are associated with lower values and brighter cells are associated with higher values. In the figure, if the communication range equals or exceeds the map side, i.e., $F \geq M$, the observed state o for all UAVs would correspond to the maps on the left and in the center. On the contrary, if $F < M$, the observation for each UAV would include a different portion of the map. It is easy to see how the swarm gains knowledge during the episode, as the drones explore the map and look for targets. In this case, the UAVs found two targets relatively quickly, and a significant portion of the grid remained unexplored.

We give reward 1 to a UAV if it is directly above a target, reward $-\theta$ if it tries to go outside the map or to position itself over an obstacle, reward $-\psi$ if it is in the same cell as another drone, and reward 0 in any other case. The UAVs will quickly learn to avoid actions that would take them outside the map or make them crash into obstacles, so the exact value of θ does not affect the final performance, but the value of ψ affects the distance that the drones try to keep from each other: if ψ is low, the drones will get close to each other if the targets are very close. Naturally, if there is a collision risk when the drones are in the same cell, the value of ψ should be high. The reward depends on \mathbf{X} , as well as on the action vector \mathbf{a} .

Indicating with \mathbf{x}_u and \mathbf{a}_u the position and action of drone u , we now define the collision variable $\chi_u(\mathbf{X}, \mathbf{A})$ as

$$\chi_u(\mathbf{X}, \mathbf{A}) = \max_{v \in (\mathcal{U} \setminus u)} \delta(\mathbf{x}_u + \mathbf{a}_u(t)\nu(\mathbf{x}_u, \mathbf{a}_u) - \mathbf{x}_v - \mathbf{a}_v(t)\nu(\mathbf{x}_v, \mathbf{a}_v)). \quad (9)$$

where $\delta(\mathbf{x})$ denotes a function that takes value 1 if the vector $\mathbf{x} = \mathbf{0}$, and zero otherwise. In short, $\chi_u(\mathbf{X}, \mathbf{A})$ has value 1 if one or more drones move to the same cell as drone u , and 0 otherwise. The collision variable depends on the moves of other agents, so the problem is distributed. The reward function for UAV u in state s if the swarm takes the joint action vector

\mathbf{A} , denoted as $r_u(s, \mathbf{A})$, is given by:

$$r_u(s, \mathbf{A}) = -\theta(1 - \nu(\mathbf{x}_u, \mathbf{a}_u)) - \psi\chi_u(\mathbf{X}, \mathbf{A}) + (1 - \chi_u(\mathbf{X}, \mathbf{A})) \sum_{k=0}^{K-1} \delta(\mathbf{x}_u + \mathbf{a}_u - \mathbf{z}_k). \quad (10)$$

In our model, the state transitions and the system observations are both deterministic; therefore, both the state evolution and the observation are not affected by random events, but only by the agents' decisions. We define a policy $\pi(\mathbf{a}_u|o_u)$ as the conditioned probability for user u to take action \mathbf{a}_u given an observation $o_u \in \mathcal{O}$. Under these assumptions, the goal of each drone u is to find the policy π^* that maximizes the cumulative expected future discounted reward $R_u(\pi) = \mathbb{E} \left[\sum_{\tau=0}^{+\infty} \gamma^\tau r_{u,\tau} | o_u, \pi \right]$, where $\gamma \in [0, 1)$ is a discount factor.

D. Distributed Deep Q-Learning

In this subsection, we will describe our Distributed Deep Q-Learning (DDQL) approach to solve the problem defined above. For the sake of readability, in the following we omit the u subscript to indicate the agent whenever possible. Each agent leverages a DQN, i.e., a Neural Network (NN) that takes as input the last observation o_t and returns the Q-values of the possible actions that can be taken, i.e., $Q(o_t, \mathbf{a})$, $\forall \mathbf{a} \in \mathcal{A}$. In Q-learning, the function $Q(o, \mathbf{a})$ is an estimate of the expected long-term reward R that will be achieved by choosing \mathbf{a} as the next action and then following the learned policy. In our case, we maintain a single DQN during the training phase, whose values are shared by all the agents. In this work, we follow the approach from [17] and leverage a *replay memory* to store the agent experience $e_t = (o_t, a_t, r_t, o_{t+1})$. Whenever the agent carries out a training step, a batch of B_{size} elements is picked from the replay memory, allowing to separate the algorithm training from the experience acquisition. The replay memory is shared between the agents during a training phase, and a new batch is used to train the agent at every step. We highlight that, in our system, all agents are the same (single DQN), and they need to generalize the problem from a limited number of states. As it would be impossible for a single UAV to experience even just a non-negligible fraction of possible states in the training, shared replay is a critical factor in the network's generalization ability. In particular, the experience replay is extremely valuable since it allows the system to improve the variety of the training samples by getting experience from the states seen by different agents. In other scenarios, it may not be convenient to exploit a shared memory, especially when the agents have to learn different tasks.

Following the DQN example from [17], we exploit the *double Q-learning* technique to remove biases from the Q-value estimation and speed up the algorithm's convergence [28]. This means that, during the training, we maintain a *target network*, whose output $Q_t(o, a)$ is used to evaluate actions, and an *update network*, whose output $Q_u(o, a)$ is used to select the policy. In particular, the bootstrap Q-value is computed as

$$Q^{\text{new}}(o_t, a_t) = r_t + \gamma \max_a Q_t(o_{t+1}, a). \quad (11)$$

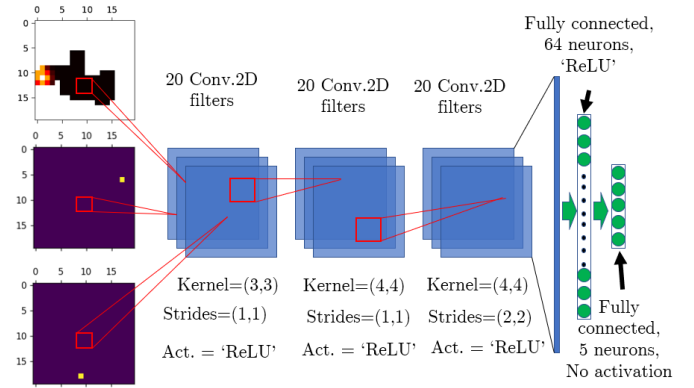


Fig. 3: Architecture of the DQN.

The value $Q^{\text{new}}(o_t, a_t)$ is then used to perform backpropagation on the update network with a learning rate set automatically by the Rectified Adam (RAdam) optimizer [29], and every n_q training steps the update network parameters are copied to the target network.

In our model, the observed state of the system for each agent can be represented by four $F \times F$ matrices, representing the agent position, the locations of the other agents, the value of explored cells, and the position of known obstacles. To simplify the state space, we consider matrices $\hat{\Phi}$ and $\hat{\Omega}$ jointly, by feeding the NN with the matrix $\hat{\Phi} - \rho\hat{\Omega}$, where ρ is a scalar parameter used to facilitate learning. Therefore, our system approximates the function $Q(o, a)$ by a Convolutional Neural Network (CNN), whose architecture is described in Fig. 3. In particular, we consider a CNN exploiting three convolutional layers followed by two fully-connected layers. The dimension of the last layer is identical to the number of actions, so that each output element can be associated to a different action $a \in \mathcal{A}$.

Hence, each agent provides training samples for the shared replay memory, which are then used in (11), so that the CNN output can converge to the Q-values $Q(o, a)$, $\forall a \in \mathcal{A}$. We implement the well-known ϵ -greedy and softmax policies to allow the agents to explore the action space during the training phase, which is carried out by simulating a sequence of episodes.

E. Computational complexity

We now discuss the computational complexity to perform one inference procedure with the neural network. We first analyze the complexity of fully-connected layers. We denote by N_k the number of neurons in the general k -th layer. To go from layer i to layer $i+1$, we need to compute the value of N_{i+1} nodes, each of which takes N_i multiplications followed by N_i additions and one non linear function, thus involving $N_{i+1}(2N_i + 1)$ operations.

We can then compute the complexity of one convolutional layer, as done in [30], when neither batch normalization nor pooling layers are present. We denote with (I_w, I_h, I_d) the shape of the input block. At layer i , we then have K_i filters with kernels dimension (W_i, H_i) , stride S_i (we use the same value along the two axes), and padding P_i . The shape of the

resulting output block will be $(\frac{I_w+2P_i-W_i}{S_i}+1, \frac{I_h+2P_i-H_i}{S_i}+1, K_i)$. The computation of each block's neuron here involves $W_i \times H_i \times I_d$ multiplications followed by the same number of additions (sum all elements plus the bias) and one non-linearity. The total number of calculations is then $(\frac{I_w+2P_i-W_i}{S_i}+1) \times (\frac{I_h+2P_i-H_i}{S_i}+1) \times K_i \times (2W_iH_iI_d+1)$.

If we consider the specific architecture of our NN reported in Fig. 3, the actual number of basic computations (multiplications, additions and non-linearities) are, respectively, 440 000, 3 704 980 and 628 180 for the three convolutional layers. The following fully-connected layers require 125 504 and 645 computations, thus the total number of operations for one decision is 4 899 309.

This computational complexity allows UAVs to take decisions in real time, as even embedded processors can deal with much more complex architectures in less than 100 ms [31]. As the physical speed of the UAVs and the much more complex vision algorithms required to identify targets are the main limiting factors for the swarm, the ND-POMDP will be performed at a relatively slow pace, with timesteps in the order of several seconds.

IV. SIMULATION SETTINGS

In this section, we describe the simulations by which we evaluated the performance of the designed system. All the results are derived through a Monte Carlo approach, where multiple independent simulations are carried out to obtain reliable statistical data. In particular, the algorithms' training is executed by carrying out a total of N_e episodes for each studied scenario (sparse or cluster), where each episode is given by N_s^t steps. Training episodes are far longer than test episodes, which have length N_s^p , since the agents need to explore the map fully.

Before training, we initialize the replay memory by executing $N_e^m = 1000$ episodes of N_s^t steps each, to allow agents to immediately start the learning procedure. If the episodes are too long, a lot of samples in which large portions of the map are already explored are added to the memory replay, and the agents will not learn properly how to move at the beginning of the episode, when the map is not explored. On the other hand, short episodes have the opposite problem, as the UAVs never learn to behave in the final parts of the episodes. A prioritized memory replay can solve this problem, but requires additional parameters. We then opted for adapting the episode length in the training phase. The even training episodes have 50 steps each, while the odd episodes have 150 steps. This alternating size prevents the replay memory from being too skewed towards situations in which the map is almost completely explored or unexplored.

Moreover, we apply transfer learning to allow the agents trained in the sparse environment to quickly adapt to the cluster scenarios (or vice-versa); to this goal, additional N_t training episodes are carried out. Finally, the performance of the proposed strategy is tested in a total of $N_p = 500$ episodes for the DDQL system. The exploration rate ε follows 2 different approaches, namely, ε -greedy and softmax. In the former, a random action is chosen with probability ε , while

the best action, i.e., the action with the highest Q-value, is chosen with probability $1-\varepsilon$. The value of ε decreases to 0 at the end of the training, since no more exploration is needed. In the latter, at each time step the probability of each action p_i is computed as the output of a softmax density function taking the Q-values as input. In this case, the temperature T decreases during the training, reducing the randomness during the selection of the actions:

$$p_i = \frac{e^{\frac{q_i}{T}}}{\sum_{j=1}^A e^{\frac{q_j}{T}}}, \quad (12)$$

where $A = 5$ is the number of actions that each drone can take. The training and testing processes are independently performed 5 times to verify the robustness of the DDQL scheme. The complete simulation settings are reported in Tab II.

To assess the performance of our DDQL scheme, we compare it with a heuristic strategy inspired by Model Predictive Control (MPC), by which drones can explore the map and reach the targets. Such a strategy is named *look-ahead* and is used as a benchmark for our analysis. The look-ahead strategy tries all possible combinations of future actions and looks at the possible future rewards, as its name suggests. In order to define it, we first define the look-ahead reward $r_u^{(\ell)}(\mathbf{X}, \mathbf{a})$ as:

$$r_u^{(\ell)}(\mathbf{X}, \mathbf{a}) = \begin{cases} \frac{\hat{\phi}(\mathbf{x}_u + \mathbf{a}_u)}{\xi(\mathbf{x}_u + \mathbf{a}_u)} & \text{if } \nu(\mathbf{x}_u, \mathbf{a}_u) = 1; \\ -\infty & \text{otherwise,} \end{cases} \quad (13)$$

where $\xi(\mathbf{x})$ is the number of UAVs located in \mathbf{x} . The look-ahead strategy never goes outside the map or on obstacles. To decide its next action, each drone u tries to maximize its expected cumulative reward over the following n_ℓ steps, assuming that none of the other drones move. In practice, the look-ahead strategy makes each drone select the action \mathbf{a}^* that maximizes

$$\max_{\mathbb{A} \in \tilde{\mathcal{A}}^{n_\ell}} \sum_{i=0}^{n_\ell-1} r_u^{(\ell)} \left(\mathbf{X} + \sum_{j=0}^{i-1} \mathbf{A}^j, \mathbf{a}^i \right), \quad (14)$$

where $\tilde{\mathcal{A}}^{n_\ell}$ is the set of ordered sequences \mathbb{A} of action vectors $\mathbb{A}^0, \mathbb{A}^1, \dots, \mathbb{A}^{n_\ell-1}$, so that $\hat{\mathbf{a}}_u^0 = \mathbf{a}^*$ and $\mathbf{a}_v^i = (0,0), \forall i \in \{0, \dots, n-1\}, v \neq u$, i.e., the set of possible move sequences of u while the other UAVs are static. If several action sequences have the same expected reward, the look-ahead strategy will choose one of them randomly. At the beginning of an episode, each drone u assumes that all the map values $\phi(\mathbf{m})$ outside its FoV are equal to 1; therefore, look-ahead forces u to continuously explore the map. However, as soon as it finds a target, u will hover over the target center. The target is then eliminated from the other agents' value maps, as it is already covered by a UAV. We highlight that the performance of look-ahead mainly depends on the n_ℓ parameter: as it increases, drones can make more foresighted decisions, but at a greater computational cost. In addition, the number of targets in the map also plays a key role in determining the computational performance: when more targets are present, we have to check whether other agents are on a target more often, in order to remove it from the map of

Parameter	Value	Description
M	{20, 24, 30, 40, 50}	Map size
F	20	Observed map size
U	{2, 3}	Number of UAVs
K	4	Number of targets
σ^2	1	Targets variance
ζ	3	Field of View
η	{0,0.1}	Obstacle frequency
d_{sparse}	8	Minimum target distance (sparse scenario)
θ	1	Obstacle/outside penalty
ψ	0.8	Collision penalty
ρ	0.2	Obstacle value
γ	0.9	Discount factor
α	Chosen by RAdam	Learning rate
N_e	{250, 750, 1000, 3000}	Training episodes
N_s^t	{50, 150}	Steps per training episode
N_s^p	40	Steps per test episode
N_t	{125, 250, 375, 750}	Transfer learning episodes
N_p	100 (LA), 500 (DDQL)	Test episodes
P_{tx}	20 dBm	Communication power
N_0	-76 dBm	Noise floor
h	40 m	UAV height
R_c	2/3	Coding rate

TABLE II: Simulation settings.

available targets. As the look-ahead strategy is computationally expensive, N_p for it was set to 100.

Finally, we also consider a scenario with a realistic communication model, in which the broadcast messages sent by each UAV at every step might be lost due to the wireless channel impairments. We used the path loss and shadowing model from [32], based on actual measurements from air-to-air communications, and considering a Rayleigh fading model with an error correction code with rate 2/3. As the simulation results will show, the physical size of the cells in the map is a critical parameter when UAVs communicate directly with each other (and not through the network infrastructure on the ground). In particular, increasing the size of the cells will impair the performance because of communication range issues: the model has an error probability of 50% at approximately 110 m, corresponding to 11 cells if a cell side is 10 m and 5 cells if the side is 20 m.

V. SIMULATION RESULTS

In what follows, we evaluate the performance of our approach in various scenarios with different characteristics.

A. Training analysis

We first consider a scenario with 2 UAVs and 4 targets in a 20×20 map. In particular, we perform multiple training phases of different duration; the longer training includes 3000 episodes, for a total of 300,000 training samples, which ensures that all our algorithms achieve convergence. The look-ahead approach is abbreviated as LA(4), as we set $n_\ell = 4$. This already had a significant computational cost, and in our simulation each look-ahead decision takes approximately 15 times longer than running a trained DDQL agent. We do not consider $n_\ell > 4$, since the computational cost of such a technique becomes excessive with limited performance gains: without coordination among the UAVs, which requires a prediction of the movements of other drones in the swarm, there is a limit on the performance of the swarm even with

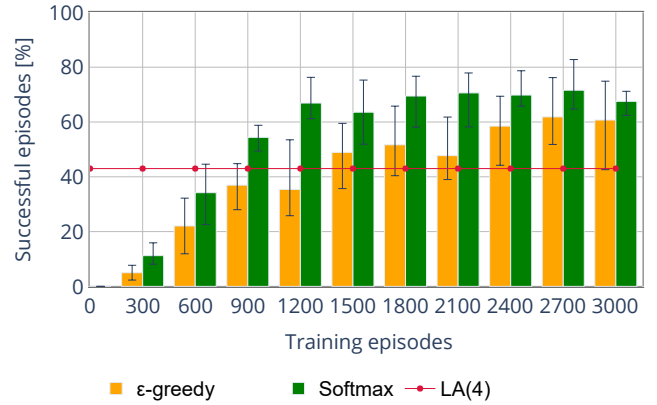


Fig. 4: Success probability over the training phase in the *cluster* scenario with 2 UAVs.

an infinite horizon. In some brief tests (which had to be on maps of a limited size due to the computational complexity of LA with a longer horizon), we noticed that LA(8) and even LA(12) show limited gains over LA(4), as the biggest factor in determining the speed at which the UAVs find the target becomes the coordination of the swarm once the horizon reaches 3 or 4 steps.

Fig. 4 shows the success probability in the *cluster* scenario as a function of the training set size and of the considered exploration profile and approach. DDQL combined with the softmax approach catches up with LA(4) in less than 900 training episodes, converging to a success probability between 0.65 and 0.7. The ϵ -greedy approach has a lower final performance and requires more time to converge with respect to the softmax profile. The error bars show the best and worst results over 5 test phases, showing that the performance improves as the UAVs gain more experience. The performance boost over the look-ahead approach is due to the DDQL scheme's ability to exploit the correlation among the target positions, quickly finding the other targets after the first one has been spotted. Instead, in the *sparse* scenario, the final performance of DDQL is similar to that of LA(4), as Fig. 5 shows. In

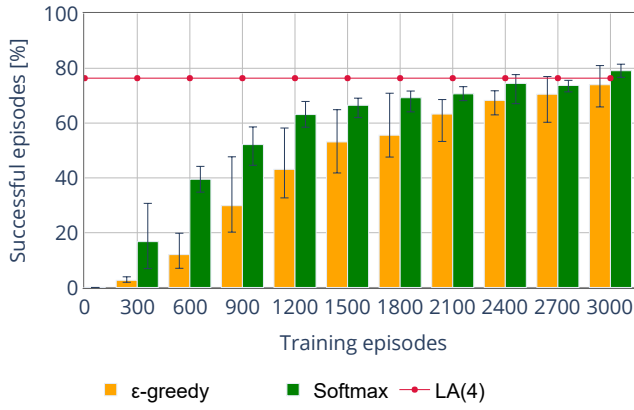


Fig. 5: Success probability over the training phase in the *sparse* scenario with 2 UAVs.

general, both DDQL and LA(4) have more success than in the cluster scenario, as finding the scattered targets is easier than finding clusters in the limited duration of an episode.

B. Success rate over time

The next set of results refer to the performance of the strategy learned by the proposed framework. Fig. 6 reports the probability of one or both drones reaching the target as a function of the number of steps. Therefore, the figure shows the trade-off between the time needed by UAVs to accomplishing their task and the success rate. In the cluster scenario (Fig. 6a), DDQL is much faster than LA, but its performance peaks out, and after 40 steps the probability of the UAVs reaching their targets does not change significantly. Indeed, we observed that, in certain cases, when a drone reaches the target, but the other one is far from any feature of the map, the latter can end up staying in place, as its Q-values for that scenario are not precise and all actions have a similar (low) value. This almost never happens before the first UAV reaches its target, since the change in the system state due to the movement of one UAV is generally enough to make the other UAV move. This is not a problem for LA, whose success rate keeps increasing with time; in the sparse scenario (Fig. 6b), LA even ends up reaching more targets than DDQL after 50 steps. The solution we found to avoid this roadblock is simply to maintain a low softmax temperature $\tau = 0.1$ even during the test phase: the bar chart shows that the DDQL Soft system is slightly slower than the greedy DDQL at the beginning, but it can avoid getting stuck. This randomization allows the agent to get out of loops, as sometimes a random sub-optimal action will change the state and allow it to reconsider, while the greedy system will keep performing the same action and remain in the same state. LA essentially does the same, randomizing its action when it is unsure which one is the best.

Fig. 7 shows one such situation: as one UAV has reached its target, while the other is far from any identified target, its Q-values will be very similar to each other, and some of the time it will stay motionless or move in small loops, as its state never changes. The fact that most of the map is still unexplored increases the probability of the UAV getting stuck, as it will

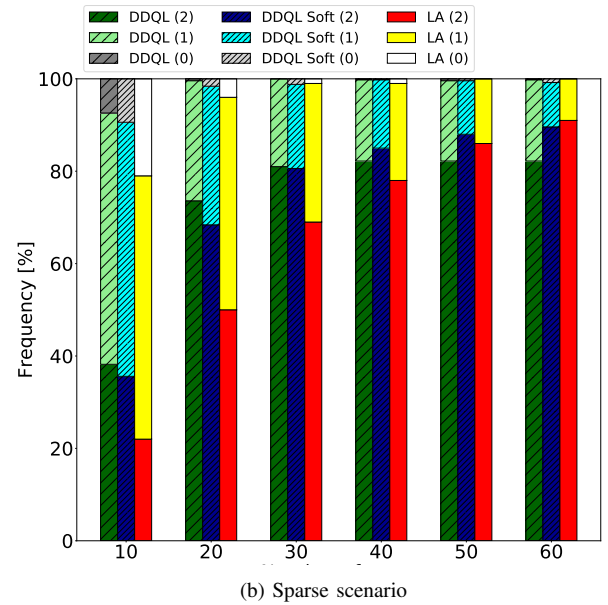
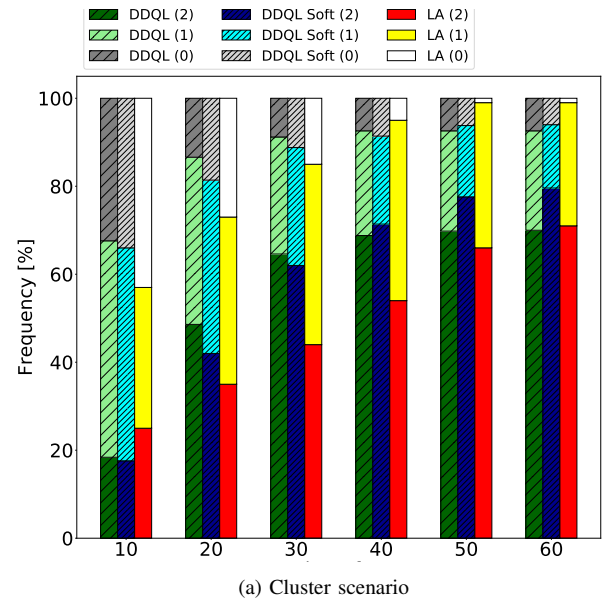


Fig. 6: The bars indicate the probability mass distribution of the number of UAVs that successfully accomplish their task (i.e., hover upon a target) by the end of the episode, when varying the duration of the episode. Each group of bars refers to the performance achieved by DDQL (with and without softmax) and by LA, in the Cluster (a) and Sparse (b) scenarios, with a total of 4 targets and 2 UAVs.

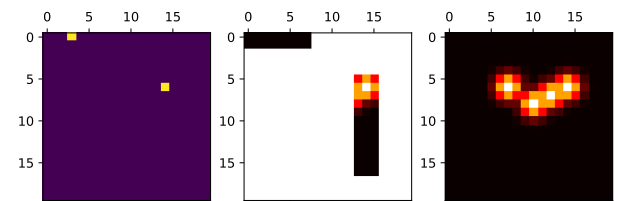


Fig. 7: Example of an episode where the second UAV is not able to reach the cluster

have limited information and its Q-values will be very similar. In the following, all the results are referred to the DDQL Soft system with $\tau = 0.1$ unless otherwise stated.

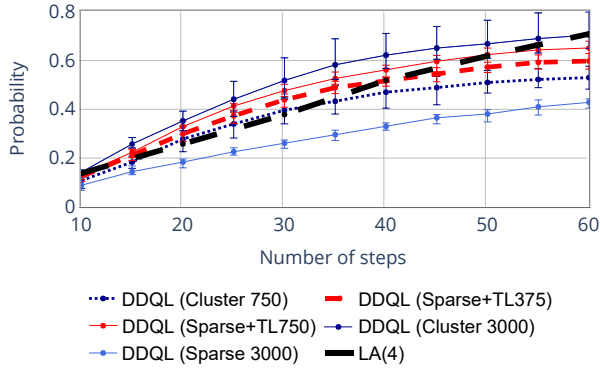


Fig. 8: CDF of the episode duration for different algorithms in the cluster scenario with 2 UAVs.

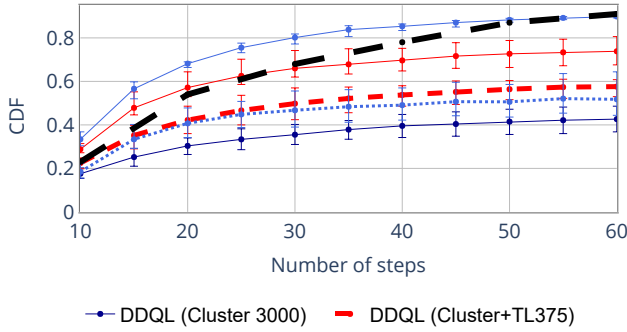


Fig. 9: CDF of the episode duration for different algorithms in the sparse scenario with 2 UAVs.

C. Adaptability and transfer learning

Here we investigate the adaptability of the proposed DDQL scheme, and the potential of the transfer learning paradigm. The latter involves the execution of an additional training phase in a different scenario than the one seen during the initial training. To this end, we consider a common target scenario, i.e., cluster (or sparse), and compare the results achieved when using strategies learned in the other domain, i.e., sparse (or cluster). More specifically, we consider the following cases:

- "Cluster N_e ": training on N_e episodes in the cluster scenario;
- "Sparse N_e ": training on N_e episodes in the sparse scenario;
- "Cluster+TL N_t ": pre-training on $N_e = 3000$ episodes in the cluster scenario, followed by an additional training of N_t episodes in the target scenario.
- "Sparse+TL N_t ": pre-training on $N_e = 3000$ episodes in the sparse scenario, followed by an additional training of N_t episodes in the target scenario.

Fig. 8 shows the Cumulative Distribution Function (CDF) of the *episode duration*, defined as the time until all the drones reach targets or the testing episode limit (here fixed to 60 steps) is reached. We also report the results for LA with four steps, LA(4), as a benchmark. Each point is hence the probability that all drones have accomplished their task by a given number of steps.

We observe that, as expected, the Cluster strategy achieves the highest success probability with a limited number of

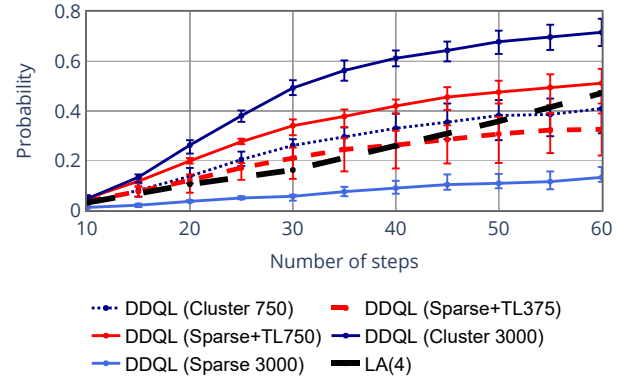


Fig. 10: CDF of the episode duration for different algorithms in the cluster scenario with 3 UAVs.

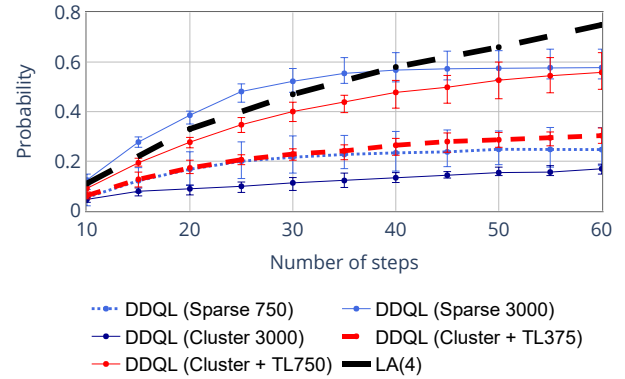


Fig. 11: CDF of the episode duration for different algorithms in the sparse scenario with 3 UAVs.

steps. LA(4) can equal its performance only when the episode duration reaches the limit of 60 steps (i.e., in less than 30% of the cases). Instead, 750 episodes of training in the cluster scenario are not sufficient to outperform LA(4), but actually enough to outperform a model trained in the sparse scenario. However, a short retraining of such model in the correct (cluster) scenario allows the algorithm to get a significant performance boost, outperforming LA(4) and getting very close to the performance of the Cluster 3000 model, which is fully trained in the correct scenario and with more than twice the number of episodes.

We repeated the experiment by swapping the role of the sparse and cluster scenarios, and changing the number of episodes during the training phase, as reflected in the legend of Fig. 9, which reports the results. As in the previous case, LA(4) meets the performance of DDQL only for episodes of 60 steps, i.e., in less than 15% of the cases. Transfer learning is again very effective, as a 750 episode re-training significantly boosts the baseline performance compared to starting from scratch. We highlight that, in general, the number of steps necessary to reach the targets is comparatively lower than in the previous scenario since, as already discussed, it is easier for UAVs to find targets in the sparse scenario.

Fig. 10 and Fig. 11 show the results for a scenario with 3 UAVs: in both cases, transfer learning is effective, but the performance is lower in the sparse scenario than in the cluster one. In this case, the risk of getting stuck is increased and

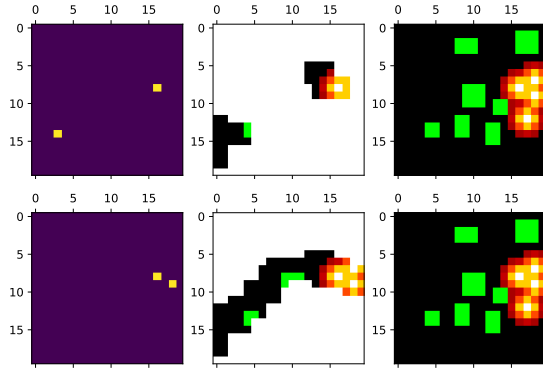


Fig. 12: Drone positions (left), known map (center), real map (right). Beginning (above) and end (below) of an episode with obstacles.

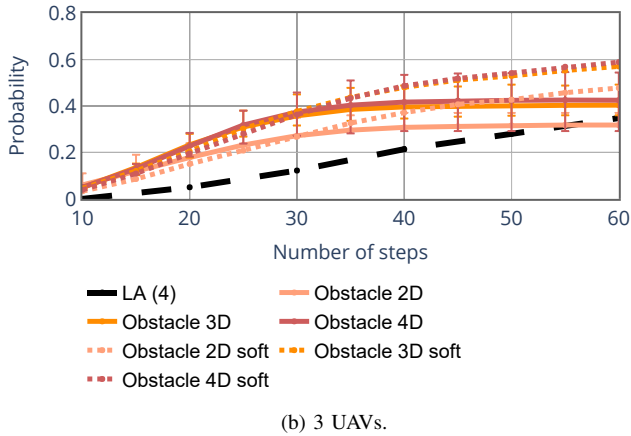
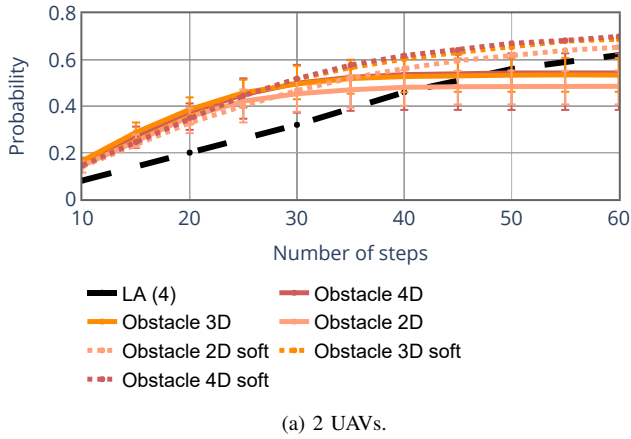
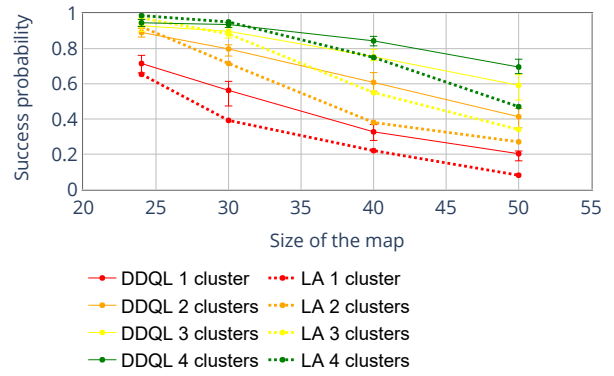


Fig. 13: CDF of the episode duration for different algorithms in the obstacle scenario.

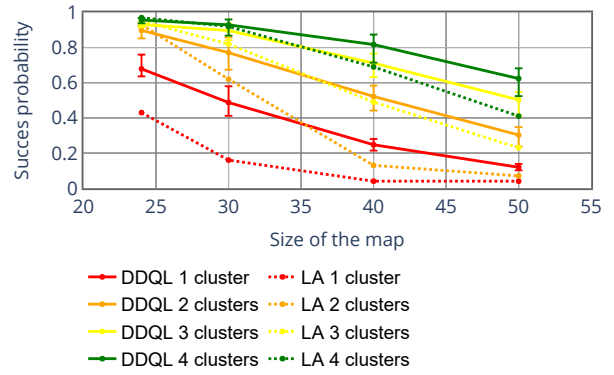
the algorithm needs more training to perform effectively in all maps.

D. Obstacles

In what follows, we consider a modified version of the cluster scenario, where some obstacles are added to the map. In particular, we empirically set the percentage of the map occupied by obstacles to 10%, searching for a balance between increased system complexity and the realism of the scenario.



(a) 2 UAVs.



(b) 3 UAVs.

Fig. 14: Success probability as a function of the map size and the number of clusters.

An example of the system state representation with obstacles is shown in Fig. 12 at the beginning and at the end of an episode. The obstacles are marked in green.

Fig. 13a shows the performance of the LA approach and DDQL in the case of 2 UAVs and 4 targets. The DDQL solution has been trained for scenarios with 2, 3 and 4 UAVs (labeled in the plots as 2D, 3D, and 4D, respectively), and then tested in the scenario with 2 and 3 UAVs, with and without the use of the softmax approach in the testing phase. In both cases, it is clear that the models trained with more UAVs are able to outperform those with fewer UAVs in both considered scenarios. Furthermore, as for the case without obstacles, the use of the softmax policy during the testing phase increases the performance, especially when the episodes are longer, as it keeps the UAVs from getting stuck. In the scenario with 3 UAVs in Fig. 13b, the performance is generally lower, meaning that the swarm needs more training. However, DDQL is able to outperform the LA approach in both cases, reaching targets significantly faster in the scenario with 3 drones.

E. Transfer learning on bigger maps with larger swarms and communication impairments

We then show how well DDQL is able to generalize to bigger maps in the testing phase. For this reason, the algorithm has been trained on a map with $M = 24$, maintaining $F = 20$, and the testing phase included bigger maps and different numbers of clusters. All the results shown in the following figures

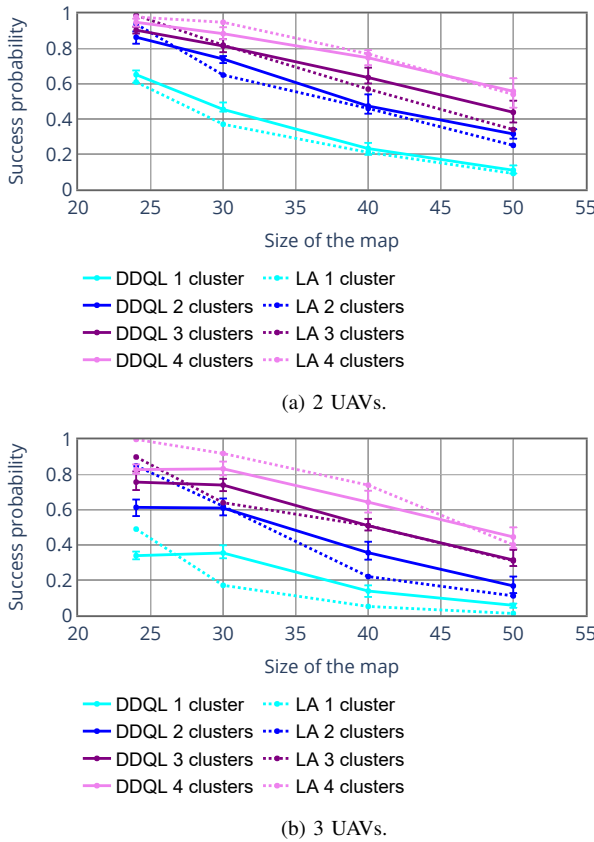


Fig. 15: Success probability as a function of the map size and the number of clusters with obstacles.

are obtained with 100-step episodes: the longer duration is needed to allow the agents to reach the targets even in bigger maps. For similar reasons, the scenarios with more clusters are studied to maintain a similar proportion of surface occupied by targets even in the bigger maps. Fig. 14a and Fig. 14b show how the performance varies as a function of the size of the environment and the number of clusters present in the map. In both cases, DDQL shows a good adaptability, getting better performance than LA in all cases, with a bigger gain in bigger maps. In Fig. 15a and Fig. 15b, the same scenarios are studied with the addition of the obstacles in the map, covering about 10% of the size of the map. In this case, DDQL will need some retraining to reach LA's performance on smaller maps, while the performance is similar when the map is bigger. However, we recall that DDQL also has a significant advantage in terms of computational cost, so it is preferable if performance is similar.

It is also interesting to test the transfer capabilities of the algorithms in more complex scenarios, including far larger swarms and imperfect communications: as DDQL relies on information from other UAVs to find targets and avoid collisions, a limited communication range can impair its performance significantly. As Fig. 16a shows, 10 drones moving in a large map with obstacles (with 16 targets in 4 clusters, as above) can coordinate effectively with no retraining, outperforming the LA approach. Performance loss is limited even with communication restrictions if each cell is a square with a 10 m side,

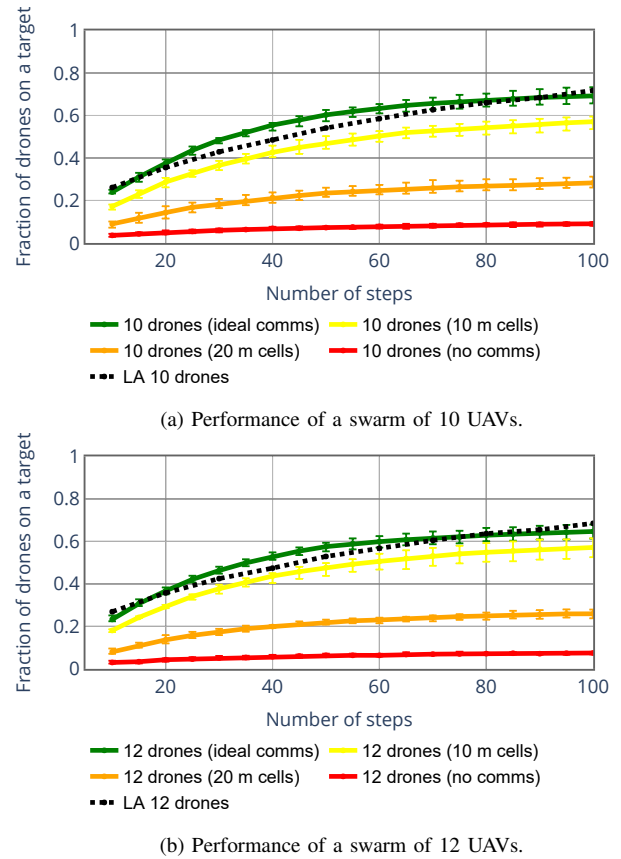


Fig. 16: Effect of imperfect communications on the performance of DDQL in a large map.

corresponding to a maximum range of about 11 cells with 50% packet loss at the boundary of the coverage area. Performance loss with respect to the perfect communication scenario is limited, confirming the intuitive idea that information from neighbors inside the visible area is the most critical to find and reach the targets. If the cell side is doubled, effectively halving the communication range and introducing significant errors even for packets between immediate neighbors, the performance drops significantly, and becomes even worse if there is no communication at all between the UAVs. This would be true for any cooperative algorithm, as information from other agents can be used to optimize the exploration of the map, but we highlight that DDQL has always been trained assuming ideal communication, and the communication impairments have been considered only in the test phase. Therefore, the UAVs might be confused by the lack of information, and a partial retraining might yield better results as the agents transfer their experience and learn to deal with the more limited feedback. On the other hand, the algorithm scales extremely well to larger swarms, slightly outperforming LA even with no retraining in the new scenario. The same pattern holds for the case with 12 drones, which is shown in Fig. 16b.

Finally, we tested an extreme transfer learning scenario, not only increasing the size of the map and the number of UAVs, but also switching from the synthetic obstacle distribution on the map to one derived from a real map. The map of obstacles was obtained from a city map of the area just east

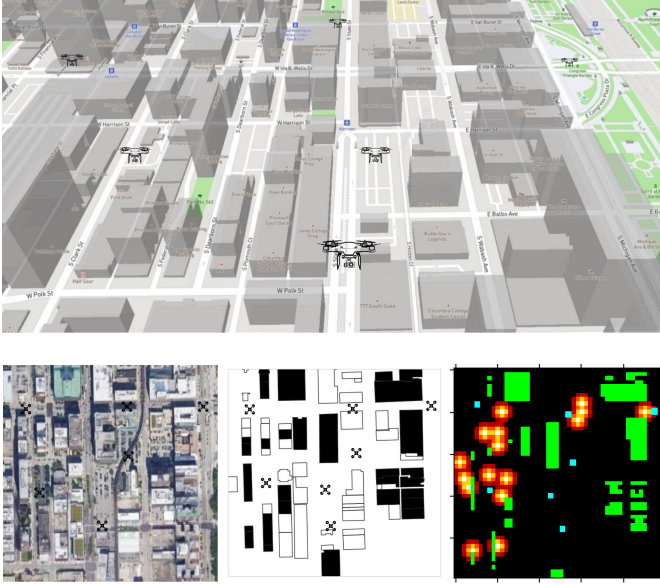


Fig. 17: Extraction of the map from building height data in a 500 m by 500 m area in the downtown Chicago Loop neighborhood.

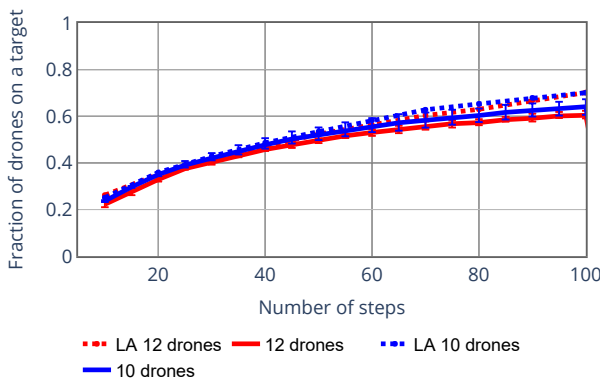


Fig. 18: Performances on the real map of Chicago

of LaSalle Street Station in downtown Chicago, in the central Loop neighborhood. As shown in Fig. 17, we obtained the height profiles of the buildings in the area, considering as obstacles their parts with a height of over 10 stories (i.e., approximately 40 m, the minimum legal hovering altitude for UAVs). The 500 m by 500 m area was then divided into 2500 square cells with 10 m sides, converting the height profile to an obstacle map in the grid. 11% of the map was occupied by obstacles, so the map was approximately as full as the one used in the training, which had 10% obstacle cover, but the individual obstacles were larger and concentrated along South Wabash Avenue and South Dearborn Street. This is an additional hurdle for DDQL, which was not trained to deal with obstacles concentrated along streets, which make it more difficult to find an appropriate path. However, as Fig. 18 shows, the DDQL system can find targets approximately as fast as LA, underperforming a little only on higher percentiles. With a modicum of retraining, DDQL should be able to adapt to the different structure, exploiting the regularities in city blocks to avoid obstacles and find targets even more quickly.

In conclusion, we have shown that DDQL is able to find efficient strategies for the UAVs to reach targets faster than look-ahead solutions in complex environments. The algorithm is scalable to larger maps, larger swarms, and limited communications without any retraining, and can deal with obstacles and very different target distributions with a limited amount of retraining. This shows that the solution is powerful and versatile, adapting easily to new conditions. However, there is still some margin for improvement, particularly in scenarios in which almost all the UAVs in the swarm have already reached a target, while the last stragglers are far from any feature in the map. This case represents most of the residual failures of the algorithm, and solving it is an important future objective.

VI. CONCLUSIONS AND FUTURE WORK

In this work, we studied the problem of area monitoring and surveillance with a swarm of drones. We modeled the environment with a 2D grid and cast the problem into the theoretical framework of ND-POMDP. We have examined various scenarios, including obstacles and maps of different sizes, and the proposed algorithm outperformed a computationally intensive look-ahead approach in almost all scenarios.

Important research directions include the introduction of dynamic targets, which would be an important step to increase the scenario's realism, as well as different roles for the drones, which can be assigned dynamically and would allow us to examine another interesting aspect of the MARL problem, increasing the difficulty of coordinating the UAVs' actions.

ACKNOWLEDGMENTS

This work has been partially supported by the U.S. Army Research Office (ARO) under Grant no. W911NF1910232, "Towards Intelligent Tactical Ad hoc Networks (TITAN)" and by MIUR (Italian Ministry for Education and Research) under the initiative "Departments of Excellence" (Law 232/2016)".

REFERENCES

- [1] N. Hossein Motlagh, T. Taleb, and O. Arouk, "Low-altitude unmanned aerial vehicles-based Internet of Things services: Comprehensive survey and future perspectives," *IEEE Internet of Things Journal*, vol. 3, no. 6, pp. 899–922, Sep. 2016.
- [2] H. Shakhathreh, 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.
- [3] R. Shakeri, M. A. Al-Garadi, A. Badawy, A. Mohamed, T. Khatat, A. K. Al-Ali, K. A. Harras, and M. Guizani, "Design challenges of multi-UAV systems in cyber-physical applications: A comprehensive survey and future directions," *IEEE Communications Surveys and Tutorials*, vol. 21, no. 4, pp. 3340–3385, 4th quarter 2019.
- [4] S. J. Chung, A. A. Paranjape, P. Dames, S. Shen, and V. Kumar, "A survey on aerial swarm robotics," *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 837–855, Aug. 2018.
- [5] D. Baldazo, J. Parras, and S. Zazo, "Decentralized multi-agent deep reinforcement learning in swarms of drones for flood monitoring," in *European Signal Processing Conference (EUSIPCO)*. EURASIP, Sep. 2019.
- [6] D. Albani, D. Nardi, and V. Trianni, "Field coverage and weed mapping by UAV swarms," in *International Conference on Intelligent Robots and Systems*. IEEE, Sep. 2017, pp. 4319–4325.
- [7] K. D. Julian and M. J. Kochenderfer, "Distributed wildfire surveillance with autonomous aircraft using deep reinforcement learning," *Journal of Guidance, Control, and Dynamics*, vol. 42, no. 8, pp. 1768–1778, Aug. 2019.

- [8] R. S. Sutton, A. G. Barto *et al.*, *Introduction to reinforcement learning*. MIT press Cambridge, Mar. 1998.
- [9] P. Hernandez-Leal, M. Kaisers, T. Baarslag, and E. M. de Cote, "A survey of learning in multiagent environments: Dealing with non-stationarity," in *3rd International Workshop on Conflict Resolution in Decision Making (COREDEMA)*, Jul. 2017.
- [10] F. Venturini, F. Mason, F. Pase, F. Chiariotti, A. Testolin, A. Zanella, and M. Zorzi, "Distributed reinforcement learning for flexible UAV swarm control with transfer learning capabilities," in *Proceedings of the 6th ACM Workshop on Micro Aerial Vehicle Networks, Systems, and Applications*, Jun. 2020.
- [11] L. Busoniu, R. Babuška, and B. De Schutter, "Multi-agent Reinforcement Learning: An Overview," *Innovations in Multi-Agent Systems and Applications*, vol. 310, pp. 113–147, Nov. 2010.
- [12] 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, Nov. 2019.
- [13] R. Zanol, F. Chiariotti, and A. Zanella, "Drone mapping through multi-agent reinforcement learning," in *Wireless Communications and Networking Conference (WCNC)*. IEEE, Apr. 2019, pp. 1–7.
- [14] J. Cui, Y. Liu, and A. Nallanathan, "The application of multi-agent reinforcement learning in UAV networks," in *International Conference on Communications Workshops (ICC)*. IEEE, May 2019.
- [15] —, "Multi-agent reinforcement learning-based resource allocation for UAV networks," *IEEE Transactions on Wireless Communications*, vol. 19, no. 2, pp. 729–743, Feb. 2020.
- [16] J. N. Foerster, Y. M. Assael, N. De Freitas, and S. Whiteson, "Learning to communicate with deep multi-agent reinforcement learning," in *Advances in Neural Information Processing Systems (NIPS)*, Dec. 2016, pp. 2145–2153.
- [17] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, Feb. 2015.
- [18] F. A. Oliehoek, C. Amato *et al.*, *A concise introduction to decentralized POMDPs*. Springer, Apr. 2016.
- [19] C. Wu, B. Ju, Y. Wu, X. Lin, N. Xiong, G. Xu, H. Li, and X. Liang, "UAV autonomous target search based on deep reinforcement learning in complex disaster scene," *IEEE Access*, vol. 7, pp. 117 227–117 245, 2019.
- [20] Y.-H. Hsu and R.-H. Gau, "Reinforcement learning-based collision avoidance and optimal trajectory planning in UAV communication networks," *IEEE Transactions on Mobile Computing*, Jun. 2020.
- [21] U. Challita, W. Saad, and C. Bettstetter, "Deep reinforcement learning for interference-aware path planning of cellular-connected UAVs," in *International Conference on Communications (ICC)*. IEEE, May 2018.
- [22] C. H. Liu, Z. Chen, J. Tang, J. Xu, and C. Piao, "Energy-efficient UAV control for effective and fair communication coverage: A deep reinforcement learning approach," *IEEE Journal on Selected Areas in Communications*, vol. 36, no. 9, pp. 2059–2070, Aug. 2018.
- [23] X. Liu, Y. Liu, and Y. Chen, "Reinforcement learning in multiple-UAV networks: Deployment and movement design," *IEEE Transactions on Vehicular Technology*, vol. 68, no. 8, pp. 8036–8049, Jun. 2019.
- [24] J. Hu, H. Zhang, L. Song, Z. Han, and H. V. Poor, "Reinforcement learning for a cellular internet of UAVs: protocol design, trajectory control, and resource management," *IEEE Wireless Communications*, vol. 27, no. 1, pp. 116–123, Feb. 2020.
- [25] R. Nair, P. Varakantham, M. Tambe, and M. Yokoo, "Networked distributed POMDPs: A synthesis of distributed constraint optimization and POMDPs," in *19th Conference on Artificial Intelligence*, vol. 5. AAAI, Jul. 2005, pp. 133–139.
- [26] A. Kumar, S. Zilberstein, and M. Toussaint, "Scalable multiagent planning using probabilistic inference," in *International Joint Conference on Artificial Intelligence (IJCAI)*. Citeseer, 2011, pp. 2140–2146.
- [27] C. Zhang and V. R. Lesser, "Coordinated Multi-Agent Reinforcement Learning in Networked Distributed POMDPs," in AAAI, Aug. 2011.
- [28] H. Van Hasselt, A. Guez, and D. Silver, "Deep reinforcement learning with Double Q-learning," in *30th Conference on Artificial Intelligence*. AAAI, Mar. 2016.
- [29] L. Liu, H. Jiang, P. He, W. Chen, X. Liu, J. Gao, and J. Han, "On the variance of the adaptive learning rate and beyond," *arXiv preprint arXiv:1908.03265*, Aug. 2019.
- [30] K. He and J. Sun, "Convolutional neural networks at constrained time cost," *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5353–5360, Jun. 2015.
- [31] S. Bianco, R. Cadene, L. Celona, and P. Napoletano, "Benchmark analysis of representative deep neural network architectures," *IEEE Access*, vol. 6, pp. 64 270–64 277, 2018.
- [32] T. Liu, Z. Zhang, H. Jiang, Y. Qian, K. Liu, J. Dang, and L. Wu, "Measurement-based characterization and modeling for low-altitude UAV air-to-air channels," *IEEE Access*, vol. 7, pp. 98 832–98 840, 2019.



Federico Venturini received the bachelor's degree in Information Engineering and the master's degree in Telecommunications Engineering from the University of Padova, in 2017 and 2019. His master's thesis project was on the application of a deep reinforcement learning framework for the coordination of a swarm of drones. He spent the 2020 as a researcher working on the extension of his master's thesis. His current research interests include Deep Reinforcement Learning and the application of Machine Learning algorithms to real world scenarios.



Federico Mason received the bachelor's and master's degree in telecommunication engineering from the University of Padua, Italy, in 2016 and 2018, respectively, where he is currently pursuing the Ph.D. degree. In 2020 he won a national grant from the GARR association to investigate new artificial intelligence solutions to enable the Network Slicing paradigm in future 5G and 6G networks. His current research interests include the analysis and development of Reinforcement Learning algorithms to optimize complex telecommunication systems.



Francesco Pase is currently a PhD student at the University of Padova, where he received his bachelor's degree in Information Engineering and his master's degree in Telecommunications Engineering, in 2017 and 2019. His master's thesis project was on graph representation, and it was developed at the EPFL and at Learn to Forecast in Switzerland, Lausanne. He spent the summer 2020 working as a research intern at InstaDeep UK, London, conducting research on maximum entropy reinforcement learning. His current research interests are in the fields of distributed and reinforcement learning, cognitive communications, wireless networks and graph analysis techniques.



Federico Chiariotti [S'15-M'19] is currently a post-doctoral researcher at the Department of Electronic Systems, Aalborg University, Denmark. He received his Ph.D. in information engineering in 2019 from the University of Padova, Italy. He received the bachelor's and master's degrees in telecommunication engineering (both *cum laude*) from the University of Padova, in 2013 and 2015, respectively. In 2017 and 2018, he was a Research Intern with Nokia Bell Labs, Dublin. He has authored over 30 published papers on wireless networks and the use of artificial intelligence techniques to improve their performance. He was a recipient of the Best Paper Award at several conferences, including the 2020 IEEE INFOCOM WCNW Workshop. His current research interests include network applications of machine learning, transport layer protocols, Smart Cities, bike sharing system optimization, and Age of Information.



Alberto Testolin received the M.Sc. degree in Computer Science and the Ph.D. degree in Psychological Sciences from the University of Padova, Italy, in 2011 and 2015, respectively. He is currently assistant professor with a joint appointment at the Department of Information Engineering and Department of General Psychology at the University of Padova. He is broadly interested in artificial intelligence, machine learning and cognitive neuroscience. His main research interests are statistical learning theory, predictive coding, sensory perception, cognitive

modeling and applications of deep learning to signal processing, networking and optimization. He is an active member of the IEEE Task Force on Deep Learning.



Andrea Zanella [S'98-M'01-SM'13] is a Full Professor at the Department of Information Engineering (DEI), University of Padova, Italy. He received the Laurea degree in Computer Engineering in 1998 from the same University and the PhD in 2001. During 2000, he spent 9 months with Prof. Mario Gerla's research team at the University of California, Los Angeles (UCLA). Andrea Zanella is one of the coordinators of the SIGnals and NETworking (SIGNET) research lab. His long-established research activities are in the fields of protocol design, optimization,

and performance evaluation of wired and wireless networks. He has been serving as Technical Area Editor for the IEEE Internet of Things Journal, and Associate Editor for the IEEE Transactions on Cognitive Communications and Networking, IEEE Communications Surveys and Tutorials, and Digital Communications and Networks.



Michele Zorzi [F'07] received his Laurea and PhD degrees in electrical engineering from the University of Padova, Italy, in 1990 and 1994, respectively. During the academic year 1992-1993 he was on leave at the University of California at San Diego (UCSD). In 1993 he joined the faculty of the Dipartimento di Elettronica e Informazione, Politecnico di Milano, Italy. After spending three years with the Center for Wireless Communications at UCSD, in 1998 he joined the School of Engineering of the University of Ferrara, Italy, where he became a professor in

2000. Since November 2003 he has been on the faculty of the Information Engineering Department at the University of Padova. His present research interests include performance evaluation in mobile communications systems, WSN and Internet of Things, cognitive communications and networking, 5G mmWave cellular systems, vehicular networks, and underwater communications and networks. He is the recipient of several awards from the IEEE Communications Society, including the Best Tutorial Paper Award (2008, 2019), the Education Award (2016), and the Stephen O. Rice Best Paper Award (2018). He was the Editor in Chief of IEEE Wireless Communications from 2003 to 2005, of the IEEE Transactions on Communications from 2008 to 2011, and of the IEEE Transactions on Cognitive Communications and Networking from 2014 to 2018. He served the IEEE Communications Society as a Member-at-Large of the Board of Governors from 2009 to 2011, as Director of Education in 2014-15, and as Director of Journals in 2020-21.