Conservative and Risk-Aware Offline Multi-Agent Reinforcement Learning

Eslam Eldeeb, Houssem Sifaou, Osvaldo Simeone, Mohammad Shehab, and Hirley Alves

Abstract—Reinforcement learning (RL) has been widely adopted for controlling and optimizing complex engineering systems such as next-generation wireless networks. An important challenge in adopting RL is the need for direct access to the physical environment. This limitation is particularly severe in multi-agent systems, for which conventional multi-agent reinforcement learning (MARL) requires a large number of coordinated online interactions with the environment during training. When only offline data is available, a direct application of online MARL schemes would generally fail due to the epistemic uncertainty entailed by the lack of exploration during training. In this work, we propose an offline MARL scheme that integrates distributional RL and conservative Q-learning to address the environment's inherent aleatoric uncertainty and the epistemic uncertainty arising from the use of offline data. We explore both independent and joint learning strategies. The proposed MARL scheme, referred to as multi-agent conservative quantile regression, addresses general risk-sensitive design criteria and is applied to the trajectory planning problem in drone networks, showcasing its advantages.

Index Terms—Offline multi-agent reinforcement learning, distributional reinforcement learning, conservative Q-learning, UAV networks
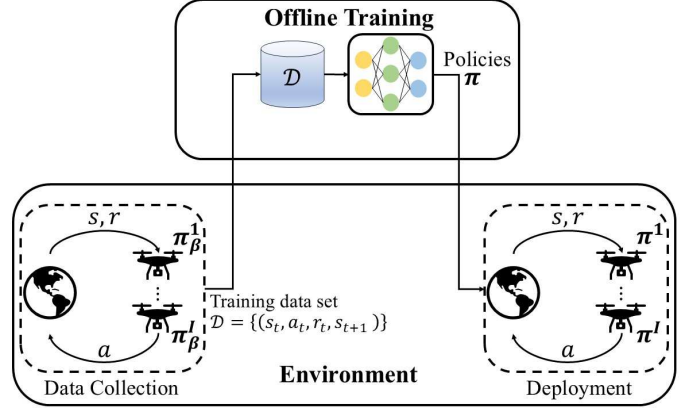
Fig. 1: Consider access to data collected offline following some fixed and unknown policies $\pi_\beta = \{\pi_\beta^i\}_{i=1}^I$ in an environment consisting of I agents. Based on this dataset, the goal is to optimize policies $\pi = \{\pi^i\}_{i=1}^I$ for the agents while ensuring robustness to the uncertainty arising from the stochastic environment, from the limited data, and from the lack of interactions with the environment.

I. INTRODUCTION

A. Context and Motivation

Recent advances in machine learning (ML) and artificial intelligence (AI), high-performance computing, cloudification, and simulation intelligence [1] have supported the development of data-driven paradigms for the engineering of complex systems, such as wireless networks [2], [3]. Reinforcement

Eslam Eldeeb and Hirley Alves are with Centre for Wireless Communications (CWC), University of Oulu, Finland. (e-mail: eslam.eldeeb@oulu.fi; hirley.alves@oulu.fi). Houssem Sifaou and Osvaldo Simeone are with the King's Communications, Learning & Information Processing (KCLIP) lab within the Centre for Intelligent Information Processing Systems (CIIPS), Department of Engineering, King's College London, WC2R 2LS London, U.K. (e-mail: houssem.sifaou@kcl.ac.uk; osvaldo.simeone@kcl.ac.uk). Mohammad Shehab is with CEMSE Division, King Abdullah University of Science and Technology (KAUST), Thuwal 23955-6900, Saudi Arabia (email: mohammad.shehab@kaust.edu.sa).

The second author has contributed to the problem definitions and the experiments. The third author has had an active role in defining the problems, as well as in writing the text, while the last two authors have had a supervisory role and have revised the text.

learning (RL) is a particularly appealing methodology for settings requiring dynamic decision-making, for which feedback can be distributed automatically or via human judgement [4], [5]. An important challenge in adopting RL solutions is their reliance on online interaction with the environment. This limitation is particularly severe in multi-agent systems, for which conventional multi-agent reinforcement (MARL) requires a large number of coordinated online interactions with the environment [6].

When only data collected offline is available, a direct application of online MARL schemes would generally fail due to the *epistemic uncertainty* entailed by the limited availability of data. In particular, even in the case of a single agent, *offline reinforcement learning*, which relies only on offline data, may over-estimate the quality of given actions that happened to perform well during data collection due to the inherent stochasticity and outliers of the environment [7]. This problem can be addressed in online RL via exploration, trying actions, and modifying return estimates based on environmental feedback. However, as mentioned, exploration is not feasible in offline RL, as policy design is based solely on the offline dataset. Furthermore, in multi-agent systems, this problem is exacerbated by the inherent uncertainty caused by the non-stationary behavior of other agents during training [8, Chapter 11].

In this paper, we propose a novel offline MARL strategy,

multi-agent conservative quantile regression (MA-CQR), that addresses the overall uncertainty caused by the use of offline data. The introduced approaches are termed *multi-agent conservative independent quantile regression* (MA-CIQR) via independent learning and *multi-agent conservative centralized quantile regression* (MA-CCQR) via joint training. These approaches integrate distributional RL [8] and offline RL [9] to support a risk-sensitive multi-agent design that mitigates impairments arising from access to limited data from the environment. We showcase the performance of MA-CIQR and MA-CCQR by focusing on the problem of designing trajectories of unmanned aerial vehicles (UAVs) used to collect data from sensors in an Internet-of-Things (IoT) scenario [10]–[12] (see Fig. 1). It is noted that this work considers the worst case in which agents can only use offline data without relying on an internal model of the environment. Future work may investigate settings in which the agents have prior information about the environment that can be used to learn a world model via methods such as model-based offline RL [13] (see Sec. I.B for a review).

B. Related Work

Offline RL: Offline RL has gained increasing interest in recent years due to its wide applicability to domains where online interaction with the environment is impossible or presents high costs and risks. Offline RL relies on a static offline transition dataset collected from the environment using some *behavioral* policy. The behavioral policy is generally suboptimal and may be unknown to the designer [7]. The reliance on a suboptimal policy for data collection distinguishes offline RL from imitation learning, in which the goal is reproducing the behavior of an expert policy [14]. The discrepancy between the behavior and optimized policies creates a *distributional shift* between training data and design objective. This shift could be resolved by collecting more data, but this is not possible in an offline setting. Therefore, the distributional shift contributes to the epistemic uncertainty of the agent.

Several approaches have been proposed to address this problem in offline RL. One class of methods constrains the difference between the learned and behavior policies [15]. Another popular approach is to learn conservative estimates of the action-value function or Q-function. Specifically, *conservative Q-learning* (CQL), proposed in [9], penalizes the values of the Q-function for *out-of-distribution* (OOD) actions. OOD actions are those whose impact is not sufficiently covered by the dataset.

Other works have leveraged offline RL with model-based RL by exploiting information about the physical environment [13], [16], [17]. The work in [13] proposed a model-based offline reinforcement learning framework that first learns the transition dynamics of the environment from the offline data and then optimizes the policy. To address the distributional shift arising in offline RL, reference [16] modified conventional model-based RL schemes by penalizing the rewards by the amount of uncertainty regarding the environment dynamics. In [17], a model-based solution for offline MARL was developed for coordination-insensitive settings. The approach

consists of learning a world model from the dataset and using it to optimize the agents' policies.

Regarding applications of offline RL to wireless systems, the recent work [18] investigated a radio resource management problem by comparing the performance of several *single-agent* offline RL algorithms.

Distributional RL: Apart from offline RL via CQL, the proposed scheme builds on *distributional RL* (DRL), which is motivated by the inherent aleatoric uncertainty caused by the stochasticity of the environment [19]–[21]. Rather than targeting the average return as in conventional RL, DRL maintains an estimate of the *distribution* of the return. This supports the design of *risk-sensitive* policies that disregard gains attained via risky behavior, favoring policies that ensure satisfactory worst-case performance levels instead.

A popular risk measure for use in DRL is the *conditional value at risk* (CVaR) [21]–[24], which evaluates the average performance by focusing only on the lower tail of the return distribution. Furthermore, a state-of-the-art DRL strategy is *quantile-regression deep Q-network* (QR-DQN), which approximates the return distribution by estimating N uniformly spaced quantiles [20].

Offline MARL: Recently, several works have been proposed that adapt the idea of conservative offline learning to the context of multi-agent systems (offline MARL) [25]–[29]. Specifically, conservative estimates of the value function in a decentralized fashion are obtained in [25] via value deviation and transition normalization. Several other works proposed centralized learning approaches. The authors in [26] leveraged first-order policy gradients to calculate conservative estimates of the agents' value functions. The work [27] presented a counterfactual conservative approach for offline MARL, while [28] introduced a framework that converts global-level value regularization into equivalent implicit local value regularization. The authors in [29] addressed the overestimation problem using implicit constraints.

Overall, all of these works focused on *risk-neutral* objectives, hence not making any provisions to address risk-sensitive criteria. In this regard, the paper [24] combined distributional RL and conservative Q-learning to develop a risk-sensitive algorithm, but only for single-agent settings.

Applications of MARL to wireless systems: Due to the multi-objective and multi-agent nature of many control and optimization problems in wireless networks, MARL has been adopted as a promising solution in recent years. For instance, related to our contribution, the work in [30] proposed an online MARL algorithm to jointly minimize the age-of-information (AoI) and the transmission power in IoT networks with traffic arrival prediction, whereas the authors in [31] leveraged MARL for AoI minimization in UAV-to-device communications. Moreover, MARL was used in [32] for resource allocation in UAV networks. The work [33] developed a MARL-based solution for optimizing power allocation dynamically in wireless systems. The authors in [34] used MARL for distributed resource management and interference mitigation in wireless networks and in [35], edge-end task division, transmit power, computing resource type matching and allocation are jointly optimized using a MARL algorithm.

Applications of Distributional RL to wireless systems: Distributional RL has been recently leveraged in [36] to carry out the optimization for a downlink multi-user communication system with a base station assisted by a reconfigurable intelligent reflector (IR). Meanwhile, reference [37] focused on the case of mmWave communications with IRs on a UAV. Distributional RL has also been used in [38] for resource management in network slicing. The paper [24] combined distributional RL and conservative Q-learning to develop a risk-sensitive algorithm, but only for single-agent settings.

All in all, to the best of our knowledge, our work in this paper is the first to integrate conservative offline RL and distributional MARL, and it is also the first to investigate the application of offline MARL to wireless systems.

C. Main Contributions

This work introduces MA-CQR, a novel offline MARL scheme that supports optimizing risk-sensitive design criteria such as CVaR. MA-CQR is evaluated on the relevant problem of UAV trajectory design for IoT networks. The contributions of this paper are summarized as follows.

- We propose MA-CQR, a novel conservative and distributional offline MARL solution. MA-CQR leverages quantile regression (QR) to support the optimization of risk-sensitive design criteria and CQL to ensure robustness to OOD actions. As a result, MA-CQR addresses both the epistemic uncertainty arising from the presence of limited data and the aleatoric uncertainty caused by the randomness of the environment.
- We present two versions of MA-CQR with different levels of coordination among the agents. In the first version, referred to as MA-CIQR, the agents' policies are optimized independently. In the second version, referred to as MA-CCQR, we leverage value decomposition techniques that allow joint training [39], [40].
- To showcase the proposed schemes, we consider a trajectory optimization problem in UAV networks [30]. As illustrated in Fig. 1, the system comprises multiple UAVs collecting information from IoT devices. The multi-objective design tackles the minimization of the AoI for data collected from the devices and the overall transmit power consumption. We specifically exploit MA-CQR to design risk-sensitive policies that avoid excessively risky trajectories in the pursuit of larger average returns.
- Numerical results demonstrate that MA-CIQR and MA-CCQR versions yield faster convergence and higher returns than the baseline algorithms. Furthermore, both schemes can avoid risky trajectories and provide the best worst-case performance. Experiments also depict that centralized training provides faster convergence and requires less offline data.

The rest of the paper is organized as follows. Section II describes the MARL setting and the design objective. Section III introduces distributional RL and conservative Q-Learning. In section IV, we present the proposed MA-CIQR algorithm using independent training, whereas section V presents the proposed MA-CCQR algorithm using centralized training. In

AoI	Age-of-information
CDF	Cumulative distribution function
CQL	Conservative Q-learning
CVaR	Conditional value at risk
DQN	Deep Q-network
DRL	Distributional reinforcement learning
MA-CCQL	Multi-agent conservative centralized Q-learning
MA-CCQR	Multi-agent conservative centralized quantile regression
MA-CIQL	Multi-agent conservative independent Q-learning
MA-CIQR	Multi-agent conservative independent quantile regression
MA-CQL	Multi-agent conservative Q-learning
MA-CQR	Multi-agent conservative quantile regression
MA-DCQN	Multi-agent deep centralized Q-network
MA-DIQN	Multi-agent deep independent Q-network
MA-DQN	Multi-agent deep Q-network
MA-QR-DCQN	Multi-agent quantile regression deep centralized Q-network
MA-QR-DIQN	Multi-agent quantile regression deep independent Q-network
MA-QR-DQN	Multi-agent quantile regression deep Q-network
MARL	Multi-agent reinforcement learning
OOD	Out-of-distribution
QR-DQN	Quantile-regression deep Q-network
UAV	Unmanned aerial vehicles

TABLE II: Notations

I	Number of agents
s_t	Overall state of the environment at time step t
a_t	Joint action of all agents at time step t
a_t^i	Action of agent i at time step t
r_t	Immediate reward at time step t
γ	Discount factor
$Q(s,a)$	Q-function
$Z(s,a)$	Return starting from (s,a)
$P(s_{t+1}|s_t,a_t)$	Transition probability
$\pi^i(a_t^i \mid s_t)$	Policy of agent i
$P_{Z(s,a)}$	Distribution of the return
$R(r_t \mid s_t, a_t)$	Stationary reward distribution
ξ	Risk tolerance level
J_ξ^{CVaR}	CVaR risk measure
$F_{Z^\pi}^{-1}(\xi)$	Inverse CDF of the return
$\mathcal{D}$	Offline dataset collected
$\theta_j^i(s,a)$	Quantile estimate of the distribution $P_{Z^i(s,a)}(\theta^i)$
$\zeta_\tau(u)$	Quantile regression Huber loss
$\Delta_{jj'}^{i(k)}$	TD errors evaluated with the quantile estimates of agent i
α	CQL hyperparameter
M	Number of devices in the system
A_t^m	AoI of device m at time step t
$g_t^{i,m}$	Channel gain between agent i and device m at time step t
$P_t^{i,m}$	Transmission power for device m to communicate with agent i at time step t
p_{risk}	Risk probability
P_{risk}	Risk penalty

Section VI, we provide numerical experiments on trajectory optimization in UAV networks. Section VII concludes the paper.

II. PROBLEM DEFINITION

In this section, we describe the multi-agent setting and formulate the problem. This discussion will be also instrumental in introducing the necessary notation, which will be leveraged in the following section to introduce important background information. We consider the setting illustrated in Fig. 1, where I agents act in a *physical environment* that evolves in discrete time as a function of the agents' actions and random dynamics. The design of the agents' policies $\pi = \{\pi^i\}_{i=1}^I$

is carried out at a central unit in Fig. 1 that has only access to a fixed dataset $\mathcal{D}$, while not being able to interact with the physical system. The dataset $\mathcal{D}$ is collected offline by allowing the agents to act in the environment according to arbitrary, fixed, and generally unknown policies $\pi_\beta = \{\pi_\beta^i\}_{i=1}^I$. In this section, we describe the multi-agent setting and formulate the offline learning problem. Tables I and II summarize the list of abbreviations and notations.

A. Multi-Agent Setting

Consider an environment characterized by a time-variant state s_t, where $t = 1, 2, \ldots$ is the discrete time index. At time step t, each agent i takes *action* $a_t^i \in \mathcal{A}^i$ within some discrete action space $\mathcal{A}^i$. We denote by $a_t = \left[a_t^1, \cdots, a_t^I\right]$ the vector of actions of all agents at timestep t. The state s_t evolves according to a *transition probability* $P(s_{t+1}|s_t, a_t)$ as a function of the current state s_t and of the action vector a_t. The transition probability $P(s_{t+1}|s_t, a_t)$ is stationary, i.e., it does not vary with time index t.

We focus on a *fully observable multi-agent reinforcement learning* setting, in which each agent i has access to the full system state s_t and produces action a_t^i by following a *policy* $\pi^i(a_t^i \mid s_t)$.

B. Design Goal

The desired goal is to find the optimal policies $\pi_*(a|s) = \{\pi_*^i(a|s)\}_{i=1}^I$ that maximize a *risk measure* $\rho(\cdot)$ of the *return* $Z^\pi = \sum_{t=0}^\infty \gamma^t r_t$, which we write as

$$J_\rho(\pi) = \rho\left[Z^\pi\right], \tag{1}$$

where $0 < \gamma < 1$ is a given discount factor. The distribution of the return Z^π depends on the policies π through the distribution of the trajectory $\mathcal{T} = (s_0, a_0, r_0, s_1, a_1, r_1, \ldots)$, which is given by

$$P(\mathcal{T}) = P(s_0)\prod_{t=0}^\infty \pi(a_t \mid s_t)R(r_t \mid s_t, a_t)P(s_{t+1} \mid s_t, a_t), \tag{2}$$

with $\pi(a_t \mid s_t) = \prod_{i=1}^I \pi^i(a_t^i \mid s_t)$ being the joint conditional distribution of the agents' actions; $P(s_0)$ being a fixed initial distribution; and $R(r_t \mid s_t, a_t)$ being the stationary *reward distribution*.

The standard choice for the risk measure in (1) is the expectation $\rho[\cdot] = \mathbb{E}[\cdot]$, yielding the standard criterion

$$J^{\text{avg}}(\pi) = \mathbb{E}\left[Z^\pi\right]. \tag{3}$$

The average criterion in (3) is considered to be *risk neutral*, as it does not directly penalize worst-case situations, catering only to the average performance.

In stochastic environments where the level of aleatoric uncertainty caused by the transition probability and/or the reward distribution is high, maximizing the expected return may not be desirable since the return Z^π has high variance. In such scenarios, designing *risk-sensitive* policies may be preferable to enhance the worst-case outcomes while reducing the average performance (3).


Fig. 2: Illustration of the conditional value-at-risk (CVaR). The quantile function $F_Z^{-1}(\xi)$ is plotted as a function of the risk tolerance level ξ. The shaded area representing the lower tail of the distribution depicts the ξ-level CVaR.

A common risk-sensitive measure is the *conditional value-at-risk* (CVaR) [22], which is defined as the conditional mean

$$J_\xi^{\text{CVaR}}(\pi) = \mathbb{E}\left[Z^\pi \mid Z^\pi \leq F_{Z^\pi}^{-1}(\xi)\right], \tag{4}$$

where $F_{Z^\pi}^{-1}(\xi)$ is the inverse cumulative distribution function (CDF) of the return Z^π for some $\xi \in [0, 1]$, i.e., the ξ-th quantile of the distribution of the return. The CVaR, illustrated in Fig. 2, focuses on the lower tail of the return distribution by neglecting values of the return that are larger than the ξ-th quantile $F_{Z^\pi}^{-1}(\xi)$. Accordingly, the probability ξ represents the *risk tolerance level*, with $\xi = 1$ recovering the risk-neutral objective (3). The CVaR can also be written as the integral of the quantile function $F_{Z^\pi}^{-1}(\xi)$ as

$$J_\xi^{\text{CVaR}}(\pi) = \frac{1}{\xi}\int_0^\xi F_{Z^\pi}^{-1}(u)\mathrm{d}u. \tag{5}$$

C. Offline Multi-Agent Reinforcement Learning

Conventional MARL [41] assumes that agents optimize their policies $\pi = \{\pi^i(a \mid s)\}_{i=1}^I$ via an online interaction with the environment, allowing for the exploration of new actions a_t as a function the state s_t. In this paper, as illustrated in Fig. 1, we assume that the design of policies is carried out on the basis solely of the availability of an offline dataset $\mathcal{D} = \{(s, a, r, s')\}$ of transitions (s, a, r, s'). Each transition follows the stationary marginal distribution from (2), with policy $\pi(a|s)$ given by the fixed and unknown *behavior policy* $\pi_\beta(a|s) = \prod_{i=1}^I \pi_\beta^i(a^i|s)$.

III. BACKGROUND

In this section, we present a brief review of distributional RL, as well as of offline RL via CQL for a single agent model [9]. This material will be useful to introduce the proposed multi-agent offline DRL solution in the next section.

A. Distributional Reinforcement Learning

Distributional RL aims at optimizing the agent's policy, π, while accounting for the inherent aleatoric uncertainty associated with the stochastic environment. To this end, it tracks the return's distribution, allowing the minimization of an arbitrary risk measure, such as the CVaR.

To elaborate, let us denote the random variable representing the return starting from a given state-action pair (s, a) as $Z^\pi(s, a)$. Taking the expectation of the return $Z^\pi(s, a)$ over distribution (2) yields the state-action value function, also known as *Q-function*, as

$$Q^\pi(s, a) = \mathbb{E}\left[Z^\pi(s, a)\right]. \tag{6}$$

Classical Q-learning algorithms learn the optimal policy π^* by finding the optimal Q-function $Q(s, a)$ as the unique fixed point of the Bellman optimality operator [42]

$$Q(s, a) = \mathbb{E}\left[r + \gamma \max_{a' \in \mathcal{A}} Q(s', a')\right], \tag{7}$$

with average evaluated with respect to the random variables $(r, s') \sim R(r|s, a)P(s'|s, a)$. The optimal policy π^* for the average criterion (3) is directly obtained from the optimal Q-function as

$$\pi^*(a|s) = \mathbb{1}\left\{a = \arg\max_{a \in \mathcal{A}} \ Q(s, a)\right\}, \tag{8}$$

with $\mathbb{1}\{\cdot\}$ being the indicator function.

Similarly, for any risk measure $\rho[\cdot]$, one can define the distributional Bellman optimality operator for the random return $Z(s, a)$ as [19], [20]

$$Z(s, a) \overset{D}{=} r + \gamma Z\left(s', \arg\max_{a' \in \mathcal{A}} \ \rho[Z(s', a')]\right), \tag{9}$$

where equality holds regarding the distribution of the random variables on the left- and right-hand sides, and the random variables (r, s') are distributed as in (7). The optimal policy for the general criterion (1) can be expressed directly as a function of the optimal $Z(s, a)$ in (9) as [19], [20]

$$\pi(a|s) = \mathbb{1}\left\{a = \arg\max_{a \in \mathcal{A}} \rho[Z(s, a)]\right\}.$$

Quantile regression DQN (QR-DQN) [20] estimates the distribution $P_{Z(s,a)}$ of the optimal return $Z(s, a)$ by approximating it via a uniform mixture of Dirac functions centered at N values $\{\theta_j(s, a)\}_{j=1}^N$, i.e.,

$$\hat{P}_{Z(s,a)}(\theta) = \frac{1}{N} \sum_{j=1}^N \delta_{\theta_j(s,a)}. \tag{10}$$

Each value $\theta_j(s, a)$ in (10) is an estimate of the quantile $F_{Z(s,a)}^{-1}(\hat{\tau}_j)$ of distribution $P_{Z(s,a)}$ corresponding to the quantile target $\hat{\tau}_j = (\tau_{j-1} + \tau_j)/2$, with $\tau_j = j/N$ for $1 \leq j \leq N$. Note that $\{\theta_j(s, a)\}_{j=1}^N$ are estimated via quantile regression, which is achieved by modeling the function mapping (s, a) to the N values $\{\theta_j(s, a)\}_{j=1}^N$ as a neural network [20], which takes a state as input, and outputs the estimated $\theta_j(s, a)$ for all actions $a \in \mathcal{A}$.

The neural network is trained by minimizing the loss

$$\frac{1}{N^2} \sum_{j=1}^N \sum_{j'=1}^N \zeta_{\hat{\tau}_j}\left(\Delta_{jj'}\right),$$

where $\Delta_{jj'}$ are *the temporal difference (TD) errors* corresponding to the quantile estimates, i.e.,

$$\Delta_{jj'} = r + \gamma \theta_{j'}(s', a') - \theta_j(s, a), \tag{11}$$

with $a' = \arg\max_{a \in \mathcal{A}} \frac{1}{N} \sum_{j'=1}^N \theta_{j'}(s', a)$, and ζ_τ is the quantile regression Huber loss defined as

$$\zeta_\tau(u) = \begin{cases} -\frac{1}{2}u^2 \left|\tau - \mathbb{1}\{u < 0\}\right|, & \text{if } |u| \leq 1 \\ \left(|u| - \frac{1}{2}\right)\left|\tau - \mathbb{1}\{u < 0\}\right|, & \text{otherwise.} \end{cases} \tag{12}$$

We refer the reader to [20] for more details about the theoretical guarantees and practical implementation of QR-DQN.

B. Conservative Q-Learning

Conservative Q-learning (CQL) is a Q-learning variant that addresses epistemic uncertainty in offline RL. Specifically, it tackles the uncertainty arising from the limited available data, which may cause some actions to be OOD due to the lack of exploration. This way, CQL is complementary to QR-DQN, which, instead, targets the inherent aleatoric uncertainty in the stochastic environment.

To introduce CQL, let us first review conventional DQN [7], which approximates the solution of the Bellman optimality condition (7) by iteratively minimizing the *Bellman loss*

$$\mathcal{L}(Q, \hat{Q}^{(k)}) = \hat{\mathbb{E}}\left[\left(r + \gamma \max_{a' \in \mathcal{A}} \hat{Q}^{(k)}(s', a') - Q(s, a)\right)^2\right], \tag{13}$$

where $\hat{\mathbb{E}}[\cdot]$ is the empirical average over samples (s, a, r, s') from the offline dataset $\mathcal{D}$; $\hat{Q}^{(k)}$ is the current estimate of the optimal Q-function Q at iteration k; and the optimization is over function $Q(s, a)$, which is typically modeled as a neural network. The term $r + \gamma \max_{a' \in \mathcal{A}} \hat{Q}^{(k)}(s', a') - Q(s, a)$ is also known as the TD-error. The only difference between offline DQN, defined in (13), and online DQNs lies in the way training data are gathered. For online DQN, the data are collected by interacting with the environment, while learning from the replay buffer, by using stochastic gradient descent (SGD). In contrast, offline DQN has access to a static offline dataset of transitions, or trajectories, that were previously generated by some unknown behavioral policy, and it learns the optimal Q-function by minimizing the Bellman error on the offline dataset over multiple epochs. In this work, we use the term "offline DQN" to refer to the basic DQN scheme designed for an offline setting. This scheme does not incorporate any modifications intended to mitigate extrapolation errors and overestimation bias that may affect an offline implementation. Considering the basic offline DQN scheme will help illustrate the failure of conventional DQN methods in offline settings in the experiments in Section VI.

The maximization over the actions in the TD error in (13) may yield over-optimistic return estimates when the Q-function is estimated using offline data. In fact, a large value of the estimated maximum return $\max_{a' \in \mathcal{A}} Q(s, a)$ may be obtained based purely on the randomness in the environment during data collection. This uncertainty could be resolved by collecting additional data. However, this is not possible in an offline setting, and hence one should consider such actions as OOD [7], [43], and count the resulting uncertainty as part of the epistemic uncertainty.

To account for this issue, the CQL algorithm adds a regularization term to the objective in (13) that penalizes excessively large deviations between the maximum estimated return $\max_{a' \in \mathcal{A}} Q(s, a)$, approximated with the differentiable quantity $\log \sum_{\tilde{a} \in \mathcal{A}} \exp(Q(s, \tilde{a}))$, and the average value of $Q(s, a)$ in the data set $\mathcal{D}$ as

$$\mathcal{L}_{\text{CQL}}(Q, \hat{Q}^{(k)}) = \frac{1}{2} \mathcal{L}(Q, \hat{Q}^{(k)}) \qquad (14)$$
$$+ \alpha \hat{\mathbb{E}} \left[\log \sum_{\tilde{a} \in \mathcal{A}} \exp(Q(s, \tilde{a})) - Q(s, a) \right],$$

where $\alpha > 0$ is a hyperparameter [9].

A combination of QR-DQN and CQL was proposed in [24] for a single-agent setting to address risk-sensitive objectives in offline learning. This approach applies a regularization term as in (14) to the distributional Bellman operator (9). The next section will introduce an extension of this approach for the multi-agent scenario under study in this paper.

IV. OFFLINE CONSERVATIVE DISTRIBUTIONAL MARL WITH INDEPENDENT TRAINING

This section proposes a novel offline conservative distributional independent Q-learning approach for MARL problems. The proposed method combines the benefits of distributional RL and CQL to address the risk-sensitive objective (1) in multi-agent systems based on offline optimization as in Fig. 1. The approaches studied here apply an *independent* Q-learning approach, whereby learning is done separately for each agent. The next section will study more sophisticated methods based on joint training.

A. Multi-Agent Conservative Independent Q-Learning

We first present a multi-agent version of CQL, referred to as *multi-agent conservative independent Q-learning* (MA-CIQL), for the offline MARL problem. As in its single-agent version described in the previous section, MA-CIQL addresses the average criterion (3), aiming to mitigate the effect of epistemic uncertainty caused by OOD actions.

To this end, each agent i maintains a separable Q-function $Q^i(s, a^i)$, which is updated at each iteration k by approximately minimizing the loss

$$\mathcal{L}_{\text{MA-CIQL}}(Q^i, \hat{Q}^{i(k)}) = \frac{1}{2} \mathcal{L}(Q^i, \hat{Q}^{i(k)}) \qquad (15)$$
$$+ \alpha \hat{\mathbb{E}} \left[\log \left(\sum_{\tilde{a}^i \in \mathcal{A}^i} \exp(Q^i(s, \tilde{a}^i)) \right) - Q^i(s, a^i) \right],$$

Algorithm 1: Conservative Independent Q-learning for Offline MARL (MA-CIQL)

Input: Discount factor γ, learning rate η, conservative penalty constant α, number of agents I, number of training iterations K, number of gradient steps G, and offline dataset $\mathcal{D}$

Output: Optimized Q-functions $Q^i(s, a^i)$ for $i = 1, ..., I$

Initialize network parameters

for *iteration k in $\{1,...,K\}$* **do**
    for *gradient step g in $\{1,...,G\}$* **do**
        Sample a batch $\mathcal{B}$ from the dataset $\mathcal{D}$
        for *agent i in $\{1,...,I\}$* **do**
            Estimate the MA-CIQL loss $\mathcal{L}_{\text{MA-CIQL}}$ in (15)
            Perform a stochastic gradient step based on the estimated loss
        end
    end
end

Return $Q^i(s, a^i) = \hat{Q}^{i(K)}(s, a^i)$ for $i = 1, ..., I$

---

which is the multi-agent version of (14) over the Q-function Q^i, where $\mathcal{L}(Q^i, \hat{Q}^{i(k)})$ is the DQN loss in (13) and $\hat{Q}^{i(k)}$ is the estimate of the Q-function of agent i at the k-th iteration. Algorithm 1 summarizes the MA-CIQL algorithm for offline MARL. Note that the algorithm applies separately to each agent and is thus an example of independent per-agent learning.

B. Multi-Agent Conservative Independent Quantile-Regression

MA-CIQL can only target the average criterion (3), thus not accounting for risk-sensitive objectives that account for the inherent stochasticity of the environment. This section introduces a risk-sensitive Q-learning algorithm for offline MARL to address the more general design objective (4) for some risk tolerance level ξ.

The proposed approach, which we refer to as *multi-agent conservative independent quantile regression* (MA-CIQR), maintains an estimate of the lower tail of the distribution of the return $Z^i(s, a)$, up to the risk tolerance level ξ, for each agent i. This is done in a manner similar to (10) by using N estimated quantiles, i.e.,

$$\hat{P}_{Z^i(s,a)}(\theta^i) = \frac{1}{N} \sum_{j=1}^{N} \delta_{\theta_j^i(s,a)}. \qquad (16)$$

Generalizing (10), however, the quantity $\theta_j^i(s, a)$ is an estimate of the quantile $F_{Z^i(s,a)}^{-1}(\hat{\tau}_j)$, with $\hat{\tau}_j = \frac{\tau_{j-1}+\tau_j}{2}$ and $\tau_j = \xi j / N$ for $1 \leq j \leq N$. This way, only the quantiles of interest cover the return distribution up to the ξ-th quantile.

At each iteration k, each agent, i, updates the distribution (16) by minimizing a loss function that combines the quantile loss used by QR-DQN and the conservative penalty

introduced by CQL. Specifically, the loss function of MA-CIQR is given by

$$\mathcal{L}_{\text{MA-CIQR}}(\theta^i, \hat{\theta}^{i(k)}) = \frac{1}{2N^2}\hat{\mathbb{E}} \sum_{j=1}^{N} \sum_{j'=1}^{N} \zeta_{\hat{\tau}_j}\left(\Delta_{jj'}^{i(k)}\right) \qquad (17)$$

$$+ \alpha\hat{\mathbb{E}}\left[\frac{1}{N}\sum_{j=1}^{N}\left[\log\sum_{\tilde{a}^i \in \mathcal{A}^i} \exp\left(\theta_j^i(s, \tilde{a}^i)\right) - \theta_j^i(s, a^i)\right]\right],$$

where $\zeta_\tau(u)$ is the quantile regression Huber loss defined in (12) and $\Delta_{jj'}^{i(k)}$ are the TD errors evaluated with the quantile estimates as

$$\Delta_{jj'}^{i(k)} = r + \gamma\hat{\theta}_{j'}^{i(k)}(s', a'^i) - \theta_j^i(s, a^i), \qquad (18)$$

where $a'^i = \arg\max_{a^i \in \mathcal{A}^i} \frac{1}{N}\sum_{j'=1}^{N} \hat{\theta}_{j'}^{i(k)}(s', a^i)$. Note that the TD error $\Delta_{jj'}^{i(k)}$ is obtained by using the j'-th quantile of the current k-th iteration to estimate the return as $r + \gamma\hat{\theta}_{j'}^{i(k)}(s', a'^i)$, while considering the j-th quantile $\theta_j^i(s, a^i)$ as the quantity to be optimized.

The corresponding optimized policy is finally obtained as

$$\pi^i(a^i|s) = \mathbb{1}\left\{a^i = \arg\max_{a^i \in \mathcal{A}^i} \frac{1}{N}\sum_{j=1}^{N} \theta_j^i(s', a^i)\right\}. \qquad (19)$$

By (5), the objective in (19) is an estimate of the CVaR at the risk tolerance level ξ. The pseudocode of the MA-CIQR algorithm is provided in Algorithm 2. As for MA-CIQL, MA-CIQR applies separately across all agents.

V. OFFLINE CONSERVATIVE DISTRIBUTIONAL MARL WITH CENTRALIZED TRAINING

The independent learning strategies studied in the previous section may fail to yield coherent policies across different agents. This section addresses this issue by introducing joint / centralized methods based on *value decomposition* [39].

A. Multi-Agent Conservative Centralized Q-Learning

With value decomposition, it is assumed that the global Q-function can be written as [39]

$$Q(s, a) = \sum_{i=1}^{I} \tilde{Q}^i(s, a^i), \qquad (20)$$

where the function $\tilde{Q}^i(s, a^i)$ indicates the contribution of the i-th agent to the overall Q-function. For conventional DQN, the Bellman loss (13) is minimized over the functions $\{\tilde{Q}^i(s, a^i)\}_{i=1}^{I}$. This problem corresponds to the minimization of the global loss

$$\mathcal{L}(\{\tilde{Q}^i\}_{i=1}^{I}, \{\hat{Q}^{i(k)}\}_{i=1}^{I}) = \hat{\mathbb{E}}\left[\left(r + \gamma\sum_{i=1}^{I}\max_{\tilde{a}^i \in \mathcal{A}^i}\hat{Q}^{i(k)}(s', a^i)\right.\right.$$

$$\left.\left. - \sum_{i=1}^{I}\tilde{Q}^i(s, a^i)\right)^2\right], \qquad (21)$$

where $\hat{Q}^{i(k)}$ is the current estimate of the contribution of agent i. In practice, every function $\tilde{Q}^i(s, a^i)$ is approximated using

---

Algorithm 2: Conservative Independent Quantile Regression for Offline MARL (MA-CIQR)

Input: Discount factor γ, learning rate η, number of quantiles N, conservative penalty constant α, number of agents I, number of training iterations K, number of gradient steps G, offline dataset $\mathcal{D}$, and CVaR parameter ξ

Output: Optimized quantile estimates $\{\theta_j^i(s, a^i)\}_{j=1}^{N}$ for all $i = 1, ..., I$

Define $\tau_i = \xi i/N$, $i = 1, ..., N$

Initialize network parameters for each agent

for *iteration k in* $\{1,...,K\}$ **do**
  for *gradient step g in* $\{1,...,G\}$ **do**
    Sample a batch $\mathcal{B}$ from the dataset $\mathcal{D}$
    for *agent i in* $\{1,...,I\}$ **do**
      for *j in* $\{1,...,N\}$ **do**
        for *j' in* $\{1,...,N\}$ **do**
          Calculate TD errors $\Delta_{jj'}^{i(k)}$ using (18)
        end
      end
    Estimate the MA-CIQR loss $\mathcal{L}_{\text{MA-CIQR}}$ in (17)
    Perform a stochastic gradient step based on the estimated loss
  end
end

Return $\{\theta_j^i(s, a^i)\}_{j=1}^{N} = \{\hat{\theta}_j^{i(K)}(s, a^i)\}_{j=1}^{N}$ for all $i = 1, ..., I$

Algorithm 3: Conservative Centralized Q-learning for Offline MARL (MA-CCQL)

Input: Discount factor γ, learning rate η, conservative penalty constant α, number of agents I, number of training iterations K, number of gradient steps G, and offline dataset $\mathcal{D}$

Output: Optimized Q-functions $Q^i(s, a^i)$ for $i = 1, ..., I$

Initialize network parameters for each agent

for *iteration k in* $\{1,...,K\}$ **do**
  for *gradient step g in* $\{1,...,G\}$ **do**
    Sample a batch $\mathcal{B}$ from the dataset $\mathcal{D}$
    Estimate the MA-CCQL loss $\mathcal{L}_{\text{MA-CCQL}}$ in (23)
    Perform a stochastic gradient step to update the network parameters of each agent
  end
end

Return $\tilde{Q}^i(s, a^i) = \hat{Q}^{i(K)}(s, a^i)$, for $i = 1, ..., I$

---

a neural network. Furthermore, the policy of each agent is obtained from the optimized function $\tilde{Q}^i(s, a^i)$ as

$$\pi^i(a^i|s) = \mathbb{1}\left\{a^i = \arg\max_{a^i \in \mathcal{A}^i} \tilde{Q}^i(s, a^i)\right\}. \qquad (22)$$

The same approach can be adopted to enhance MA-CIQL

Algorithm 4: Conservative Centralized Quantile Regression for Offline MARL (MA-CCQR)

Input: Discount factor γ, learning rate η, number of quantiles N, conservative penalty constant α, number of agents I, number of training iterations K, number of gradient steps G, offline dataset $\mathcal{D}$, and CVaR parameter ξ

Output: Optimized functions $\{\tilde{\theta}^i_j(s,a^i)\}_{j=1}^N$ for all $i=1,...,I$

Define $\tau_i = \xi i/N, \ i = 1,...,N$

Initialize network parameters for each agent

for *iteration k in $\{1,...,K\}$* **do**
  for *gradient step g in $\{1,...,G\}$* **do**
    Sample a batch $\mathcal{B}$ from the dataset $\mathcal{D}$
    for *j in $\{1,...,N\}$* **do**
      for *j' in $\{1,...,N\}$* **do**
        Calculate global TD error $\Delta^{(k)}_{jj'}$ using (26)
      end
    end
    Estimate the MA-CCQR loss $\mathcal{L}_{\text{MA-CCQR}}$ in (25)
    Perform a stochastic gradient step to update the network parameters of each agent
  end
end
Return $\{\hat{\theta}^{i(K)}_j(s,a^i)\}_{j=1}^N$, for $i=1,...,I$
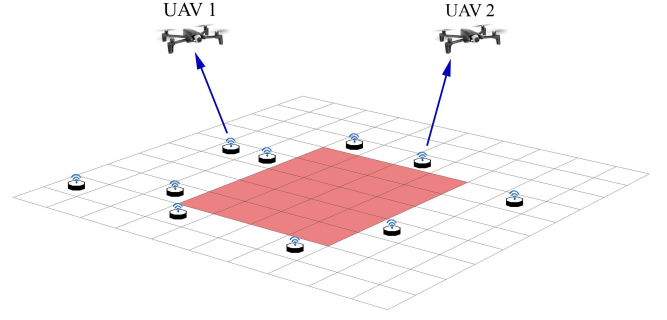

Fig. 3: Multiple UAVs serve limited-power sensors to minimize power expenditure while also minimizing the age of information for data retrieval from the sensors. The environment is characterized by a risk region for navigation of the UAVs in the middle of the grid world (colored area).

by using (20) in the loss (15). This yields the loss

$$\mathcal{L}_{\text{MA-CCQL}}(\{\tilde{Q}^i\}_{i=1}^I, \{\hat{Q}^i\}_{i=1}^I) = \frac{1}{2}\mathcal{L}(\{\tilde{Q}^i\}_{i=1}^I, \{\hat{Q}^i\}_{i=1}^I)$$
$$+ \alpha\hat{\mathbb{E}}\sum_{i=1}^I\left[\log\left(\sum_{\tilde{a}^i\in\mathcal{A}^i}\exp(\tilde{Q}^i(s,\tilde{a}^i))\right) - \tilde{Q}^i(s,a^i)\right], \quad (23)$$

with $\mathcal{L}(\{\tilde{Q}^i\}_{i=1}^I, \{\hat{Q}^i\}_{i=1}^I)$ defined in (21). The obtained scheme, whose steps are detailed in Algorithm 3, is referred to as *multi-agent conservative centralized Q-learning* (MA-CCQL). The optimized policy is given in (22).

B. Multi-Agent Conservative Centralized Quantile-Regression

The joint training approach based on the value decomposition (20) can also be applied to MA-CIQR to obtain a centralized training version referred to as *multi-agent conservative centralized quantile regression* (MA-CCQR).

To this end, we first recall that the lower tail of the distribution of $Z(s,a)$ is approximated by MA-CIQR as in (16) using the estimates of the quantiles $F_{Z(s,a)}^{-1}(\hat{\tau}_j)$, with $\hat{\tau}_j = \frac{\tau_{j-1}+\tau_j}{2}$ and $\tau_j = \xi j/N$ for $1 \le j \le N$. To jointly optimize the agents' policies, we decompose each quantile $\theta_j(s,a)$ as

$$\theta_j(s,a) = \sum_{i=1}^I \tilde{\theta}^i_j(s,a^i), \quad (24)$$

where $\tilde{\theta}^i_j(s,a^i)$ represents the contribution of agent i. The functions $\{\{\tilde{\theta}^i_j(s,a^i)\}_{i=1}^I\}_{j=1}^N$ are jointly optimized using a

loss obtained by plugging the decomposition (24) into (17) to obtain

$$\mathcal{L}_{\text{MA-CCQR}}(\{\{\tilde{\theta}^i_j\}_{i=1}^I\}_{j=1}^N, \{\{\hat{\theta}^{i(k)}_j\}_{i=1}^I\}_{j=1}^N) =$$
$$\frac{1}{2N^2}\hat{\mathbb{E}}\sum_{j=1}^N\sum_{j'=1}^N \zeta_{\hat{\tau}_j}\left(\Delta^{(k)}_{jj'}\right) \quad (25)$$
$$+ \alpha\hat{\mathbb{E}}\sum_{i=1}^I\left[\frac{1}{N}\sum_{j=1}^N\left[\log\sum_{\tilde{a}\in\mathcal{A}}\exp\left(\tilde{\theta}^i_j(s,\tilde{a}^i)\right) - \tilde{\theta}^i_j(s,a^i)\right]\right],$$

where $\{\{\hat{\theta}^{i(k)}_j\}_{i=1}^I\}_{j=1}^N$ represents the current estimate of the contribution of agent i and $\Delta^{(k)}_{jj'}$ is given by

$$\Delta^{(k)}_{jj'} = r + \gamma\sum_{i=1}^I\hat{\theta}^{i(k)}_{j'}(s',a'^i) - \sum_{i=1}^I\tilde{\theta}^i_j(s,a^i), \quad (26)$$

with $a'^i = \arg\max_{a^i\in\mathcal{A}^i}\frac{1}{N}\sum_{j'=1}^N\hat{\theta}^{i(k)}_{j'}(s',a^i)$. The individual policies of the agent are finally obtained as

$$\pi^i(a^i|s) = \mathbb{1}\left\{a^i = \arg\max_{a^i\in\mathcal{A}^i}\ \frac{1}{N}\sum_{j=1}^N\tilde{\theta}^i_j(s,a^i)\right\}.$$

For each agent, the function that maps (s,a^i) to the N values $\{\tilde{\theta}^i_j(s,a)\}_{j=1}^N$ is modeled as a neural network and the steps of the MA-CCQR scheme are provided in Algorithm 4.

VI. APPLICATION: TRAJECTORY LEARNING IN UAV NETWORKS

In this section, we consider the application of offline MARL to the trajectory optimization problem in UAV networks. Following [44], as illustrated in Fig. 3, we consider multiple UAVs acting as BSs to receive uplink updates from limited-power sensors.

A. Problem Definition and Performance Metrics

Consider a grid world, as shown in Fig. 3, where each cell is a square of length L_c. The system comprises a set $\mathcal{M}$ of M uplink IoT devices deployed uniformly in the grid world. The devices report their observations to I fixed-velocity rotary-wing UAVs flying at height h and starting from positions

selected randomly on the grid. The grid world contains normal cells, represented as white squares, and a *risk region* of special cells, colored in the figure. The risk region can be an area with a high probability of UAV collision and/or locations with a high chance of signal blockages. The current position at each time t of each UAV i is projected on the plane as coordinates (x_t^i, y_t^i). The goal is to determine trajectories for the UAVs on the grid that jointly minimize the AoI and the transmission powers across all the IoT devices.

The AoI measures the freshness of the information collected by the UAVs from the devices [44]. For each device m, the AoI is defined as the time elapsed since the last time data from the device was collected by a UAV [45], [46]. Accordingly, the AoI of device m at time t is updated as follows

$$A_t^m = \begin{cases} 1, & \text{if } V_t^m = 1, \\ \min\{A_{\max}, A_{t-1}^m + 1\}, & \text{otherwise;} \end{cases} \quad (27)$$

where $A_{\max}$ is the maximum AoI, and $V_t^m = 1$ indicates that device m is served by a UAV at time step t. The maximum value $A_{\max}$ determines the maximum penalty assigned to the UAVs for not collecting data from a device at any given time.

For the sake of demonstrating the idea, we assume line-of-sight (LoS) communication links and write the channel gain between agent i and device m at time step t as

$$g_t^{i,m} = \frac{g_0}{h^2 + (L_t^{i,m})^2}, \quad (28)$$

where g_0 is the channel gain at a reference distance of 1 m and $L_t^{i,m}$ is the distance between UAV i and device m at time t. Using the standard Shannon capacity formula, for device m to communicate to UAV i at time step t, the transmission power must be set to [47]

$$P_t^{i,m} = \frac{\left(2^{\frac{E}{B}} - 1\right)\sigma^2}{g_t^{i,m}}, \quad (29)$$

where E is the size of the transmitted packet, B is the bandwidth, and σ^2 is the noise power.

If all the UAVs are outside the risk region, the reward function is given deterministically as a weighted combination of the sums of AoI and powers across all agents

$$r_t = -\frac{1}{M}\sum_{m=1}^{M} A_t^m - \lambda \sum_{m=1}^{M} P_t^{i_m, m}, \quad (30)$$

where $\lambda > 0$ is a parameter that controls the desired trade-off between AoI and power consumption. In contrast, if any of the UAVs is within the risk region, with probability p_{risk}, the reward is given by (30) with the addition of a penalty value $\text{P}_{\text{risk}} > 0$, while it is equal to (30) otherwise. For instance, if the risk region corresponds to an area with a high probability of signal blockages, the penalty P_{risk} may be chosen to be proportional to the amount of power needed to resend the packets lost due to blockages of the communication links between the UAVs and the sensors. That said, it is emphasized that the proposed model is general and that the specific application scenario would practically dictate the choice of the penalty P_{risk}.

TABLE III: Simulation parameters and hyperparameters

Parameter	Value	Parameter	Value
g_0	30 dB	α	1
B	1 MHz	γ	0.99
h	100 m	ξ	0.15
E	5 Mb	σ^2	-100 dBm
γ	0.99	$A_{\max}$	100
Batch size	128	λ	500
Iterations K	150	p_{risk}	0.1
L_c	100 m	Optimizer	Adam

To complete the setting description, we define state and actions as follows. The global state of the system at each time step t is the collection of the UAVs' positions and the individual AoI of the devices, i.e., $s_t = [x_t^1, y_t^1, \cdots, x_t^I, y_t^I, A_t^1, A_t^2, \cdots, A_t^M]$. At each time t, the action $a_t^i = [w_t^i, d_t^i]$ of each UAV i includes the direction $w_t^i \in \{\text{north}, \text{south}, \text{east}, \text{west}, \text{hover}\}$, where "hover" represents the decision of staying in the same cell, while the other actions move the UAV by one cell in the given direction. It also includes the identity $d_t^i \in \mathcal{M} \cup \{0\}$ of the device served at time t, with $d_t^i = 0$ indicating that no device is served by UAV i.

B. Implementation and Dataset Collection

We consider a 10×10 grid world with $I = 2$ UAVs serving $M = 10$ limited-power sensors and a 5×4 risk region in the middle of the grid world as illustrated in Fig. 3. We use a fully connected neural network with two hidden layers of size 256 and ReLU activation functions to represent the Q-function and the quantiles. The experiments are implemented using Pytorch on a single NVIDIA Tesla V100 GPU. Table III shows the UAV network parameters and the proposed schemes' hyperparameters. We compare the proposed method MA-CQR to baseline offline MARL schemes, namely multi-agent deep Q-network (MA-DQN) [48], MA-CQL (see Sec.IV-A), and multi-agent quantile regression DQN (MA-QR-DQN). MA-DQN corresponds to MA-CQL when no conservative penalty for OOD is applied, i.e., $\alpha = 0$ in (15), whereas MA-QR-DQN corresponds to MA-CQR when $\alpha = 0$ and $\xi = 1$. Both the independent and centralized training frameworks apply to MA-DQN, yielding multi-agent deep independent Q-network (MA-DIQN) and multi-agent deep centralized Q-network (MA-DCQN), and to MA-QR-DQN, yielding multi-agent quantile regression deep independent Q-network (MA-QR-DIQN) and multi-agent quantile regression deep centralized Q-network (MA-QR-DCQN).

For the proposed MA-CQR, we consider two settings for the risk tolerance level ξ, namely $\xi = 1$ and $\xi = 0.15$, with the former corresponding to a risk-neutral design. We refer to the former as MA-CQR and the latter as MA-CQR-CVaR. For all distributional RL schemes (MA-QR-DQN and MA-CQR in all its variants), the learning rate is set to 10^{-5}, while for all other schemes, we use a learning rate of 10^{-4}.

The offline dataset $\mathcal{D}$ is collected using online independent DQN agents. In particular, we train the UAVs using an online MA-DQN algorithm until convergence and use 6% and 16%

Fig. 4: Average test return as a function of the number of training epochs using $P_{risk} = 100$ and 16% offline dataset for a system of 2 UAVs serving 10 sensors. The return is averaged over 100 test episodes at the end of each training epoch and shown upon division by 1000.

of the total number of transitions from the observed experience as the offline datasets[1].

C. Numerical Results

First, we show the simulation results of the proposed model via independent Q-learning compared to the baseline schemes. Then, we investigate the benefits of the joint training approach compared to independent training.

1) Independent Learning: Fig. 4 shows the average test return, evaluated online using 100 test episodes, for the policies obtained after a given number of training epochs. The figure thus reports the actual return obtained by the system as a function of the computational load, which increases with the number of training epochs.

We first observe that both MA-DIQN and MA-QR-DIQN, designed for online learning, fail to converge in the offline setting at hand. This well-known problem arises from over-estimating Q-values corresponding to OOD actions in the offline dataset [7]. In contrast, conservative strategies designed for offline learning, namely MA-CIQL, MA-CIQR, and MA-CIQR-CVaR, exhibit an increasing average return as a function of the training epochs. In particular, the proposed MA-CIQR and MA-CIQR-CVaR provide the fastest convergence, needing around 30 training epochs to reach the maximum return. In contrast, MA-CIQL shows slower convergence. This highlights the benefits of distributional RL in handling the inherent uncertainties arising in multi-agent systems from the environment and the actions of other agents [8].

In Fig 5, we report the optimal achievable trade-off between sum-AoI and sum-power consumption across the devices. This region is obtained by training the different schemes while sweeping the hyperparameter values λ. We recall that
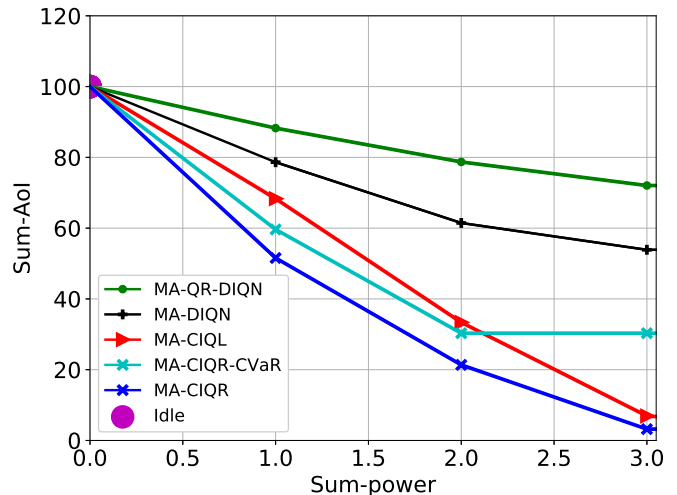
Fig. 5: Sum-AoI as a function of the sum-power using $P_{risk} = \lambda/4$ and 16% offline dataset for a system of 2 UAVs serving 10 sensors.
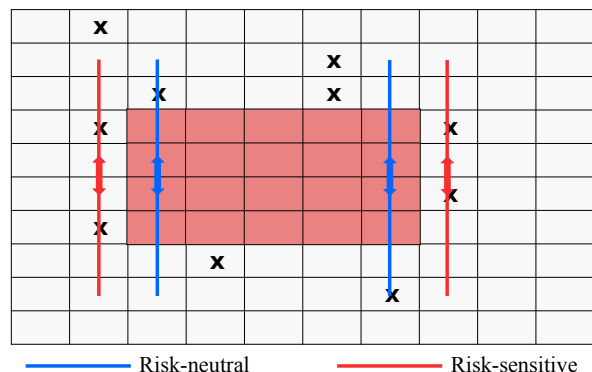

Fig. 6: A comparison between the trajectories of two UAVs using risk-neutral and risk-sensitive policies obtained via MA-CIQR and MA-CIQR-CVaR, respectively. Crosses represent the positions of the devices.

the hyperparameter λ controls the weight of the power as compared to the AoI in the reward function (30). In particular, setting $\lambda \to 0$ minimizes the AoI only, resulting in a round-robin optimal policy. At the other extreme, setting a large value of λ causes the UAV never to probe the devices, achieving the minimum power equal to zero, and the maximum AoI $A_{max} = 100$. This point is denoted as "idle point" the figure. The other curves represent the minimum sum-AoI achievable as a function of the sum-power.

From Fig. 5, we observe that the proposed MA-CIQR always achieves the best age-power trade-off with the least age and sum-power consumption within all the curves. As in Fig. 4, MA-DIQN and MA-QR-DIQN provide the worst performance due to their failure to handle the uncertainty arising from OOD actions.

In the next experiment, we investigate the capacity of the proposed risk-sensitive scheme MA-CIQR-CVaR to avoid risky trajectories. As a first illustration of this aspect, Fig. 6 shows two examples of trajectories obtained via MA-CIQR

and MA-CIQR-CVaR. It is observed that the risk-neutral policies obtained by MA-CIQR take shortcuts through the risk area, while the risk-sensitive trajectories obtained via MA-CIQR-CVaR avoid entering the risk area.

2) Centralized Learning: Here, we compare the centralized training approach with independent learning. Fig. 7 shows the average test return as a function of training epochs for an environment of 2 agents and 10 sensors. We use two offline datasets with different sizes, equal to 6% and 16% of the total transitions from the observed experience of online DQN agents. We increase the value of the risk region penalty to $P_{risk} = 300$ compared to the previous subsection experiments.

Fig. 7a elucidates that the performance of the independent learning schemes is affected by increasing the risk penalty P_{risk}. Specifically, MA-CIQL fails to reach convergence, while MA-CIQR and MA-CIQR-CVaR reach near-optimal performance but with a slower and less stable convergence than their centralized variants. However, as in Fig. 7b, a significant performance gap is observed between the proposed independent schemes and their centralized counterpart for reduced dataset size. This joint training approach in coordinating between agents during training, requiring less data to obtain effective policies. Finally, we note the joint training approach did not enhance the performance of MA-DCQN and MA-QR-DCQN as these schemes are still heavily affected by the distributional shift in the offline setting.

In a manner similar to Fig. 5, Fig. 8 shows the trade-off between sum-AoI and sum-power consumption for both independent and centralized training approaches for a system of 3 agents serving 15 sensors. Increasing the parameter λ in (30) reduces the total power consumption at the expense of AoI. Here again, we observe a significant gain in performance for MA-CCQR and MA-CCQR-CVaR compared to their independent variants. In contrast, the non-distributional schemes, MA-CIQL and MA-CCQL, show similar results, as both perform poorly in this low data regime. We also observe that the average return performance of MA-CIQR-CVaR is worse than that of MA-CIQR, while MA-CCQR-CVaR provides a comparable performance as its risk-neutral counterpart MA-CCQR. This result suggests that, while producing low-risk trajectories, MA-CIQR-CVaR can yield lower average returns as compared to the risky trajectories of its risk-neutral counterpart, MA-CIQR. In contrast, thanks to the higher level of coordination between the agents in the joint training approach, MA-CCQR-CVaR can find low-risk trajectories while maintaining a comparable average return as MA-CCQR.

Finally, to gain further insights into the comparison between MA-CCQR and MA-CCQR-CVaR, we leverage two metrics as in [24], namely the *percentage of violations* and the $CVaR_{0.15}$ return. The former is the percentage of timesteps at which one of the UAVs enters the risk region with respect to the total number of timesteps. In contrast, the $CVaR_{0.15}$ metric is the average return of the 15% worst episodes.

In Table IV, we report these two metrics, as well as the average return, with all returns normalized by 1000. Thanks to the ability of MA-CCQR-CVaR to learn how to avoid the risk region, this scheme has the lowest percentage of violations
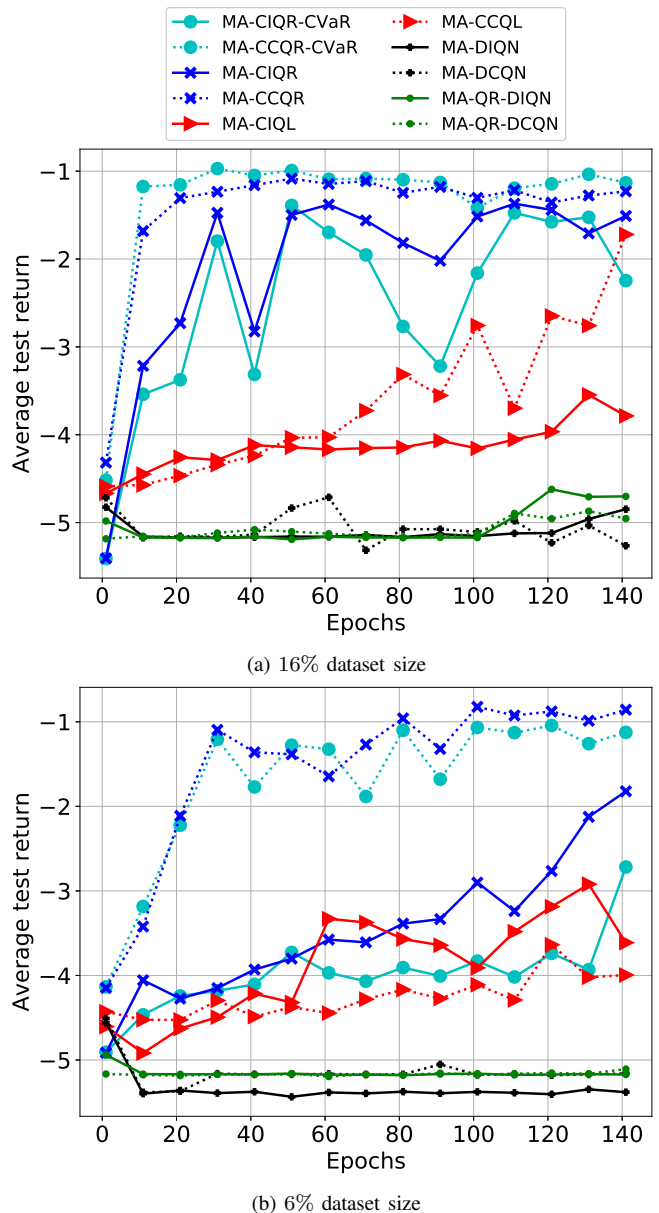

(a) 16% dataset size

(b) 6% dataset size

Fig. 7: Average test return as a function of the number of training epochs for penalty $P_{risk} = 300$ in a system of 3 UAVs and 15 sensors with data set size equal to (a) 16% and (b) 6%. The return is averaged over 100 test episodes at the end of each training epoch, and shown upon division by 1000.

TABLE IV: Performance evaluation over 100 test episodes after 150 training iterations for penalty $P_{risk} = \lambda/2.5$ in a system of 3 UAVs and 15 sensors (data set size equal to 6%).

Algorithm	Average return	$CVaR_{0.15}$ return	Violations
MA-DQN (online)	-1.5633	-1.9611	11.83%
MA-DCQN	-4.8993	-5.4930	8.29%
MA-QR-DCQN	-4.2987	-4.5743	6.85%
MA-CCQL	-3.8518	-4.4695	17.7%
MA-CCQR	-1.4028	-2.1775	8.56%
MA-CCQR-CVaR	$\mathbf{-1.3641}$	$\mathbf{-1.8513}$	**5.83%**

among all the schemes. In addition, it achieves the largest $CVaR_{0.15}$ return, with a small gain in terms of average return as compared to MA-CCQR. This demonstrates the advantages of the risk-sensitive design of policies.

Fig. 8: Sum-AoI as a function of the sum-power for penalty $P_{\text{risk}} = \lambda/2.5$ in a system of 3 UAVs and 15 sensors (data set size equal to 6%).

VII. Conclusions

In this paper, we developed a distributional and conservative offline MARL scheme for wireless systems. We considered optimizing the CVaR of the cumulative return to obtain risk-sensitive policies. We introduced two variants of the proposed scheme depending on the level of coordination between the agents during training. The proposed algorithms were applied to the trajectory optimization problem in UAV networks. Numerical results illustrate that the learned policies avoid risky trajectories more effectively and yield the best performance compared to the baseline MARL schemes. The proposed approach can be extended by considering online fine-tuning of the policies in the environment to handle the possible changes in the deployment environment compared to the one generating the offline dataset. Finally, the analysis of model-based offline MARL, which can leverage information about the physical system to learn a model of the environment, is left for future work.

References

[1] A. Lavin, D. Krakauer, H. Zenil, J. Gottschlich, T. Mattson, J. Brehmer, A. Anandkumar, S. Choudry, K. Rocki, A. G. Baydin *et al.*, "Simulation intelligence: Towards a new generation of scientific methods," *arXiv preprint arXiv:2112.03235*, 2021.

[2] C.-X. Wang, M. Di Renzo, S. Stanczak, S. Wang, and E. G. Larsson, "Artificial intelligence enabled wireless networking for 5G and beyond: Recent advances and future challenges," *IEEE Wireless Communications*, vol. 27, no. 1, pp. 16–23, 2020.

[3] M. Chen, U. Challita, W. Saad, C. Yin, and M. Debbah, "Artificial neural networks-based machine learning for wireless networks: A tutorial," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 4, pp. 3039–3071, 2019.

[4] N. C. Luong, D. T. Hoang, S. Gong, D. Niyato, P. Wang, Y.-C. Liang, and D. I. Kim, "Applications of deep reinforcement learning in communications and networking: A survey," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 4, pp. 3133–3174, 2019.

[5] W. Chen, X. Qiu, T. Cai, H.-N. Dai, Z. Zheng, and Y. Zhang, "Deep reinforcement learning for internet of things: A comprehensive survey," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 3, pp. 1659–1692, 2021.

[6] S. V. Albrecht, F. Christianos, and L. Schäfer, *Multi-Agent Reinforcement Learning: Foundations and Modern Approaches*. MIT Press, 2024. [Online]. Available: https://www.marl-book.com

[7] S. Levine, A. Kumar, G. Tucker, and J. Fu, "Offline reinforcement learning: Tutorial, review, and perspectives on open problems," *arXiv preprint arXiv:2005.01643*, 2020.

[8] M. G. Bellemare, W. Dabney, and M. Rowland, *Distributional Reinforcement Learning*. MIT Press, 2023, http://www.distributional-rl.org.

[9] A. Kumar, A. Zhou, G. Tucker, and S. Levine, "Conservative Q-learning for offline reinforcement learning," in *Advances in Neural Information Processing Systems*, vol. 33. Curran Associates, Inc., 2020, pp. 1179–1191. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2020/file/0d2b2061826a5df3221116a50

[10] M. A. Abd-Elmagid, A. Ferdowsi, H. S. Dhillon, and W. Saad, "Deep reinforcement learning for minimizing age-of-information in UAV-assisted networks," in *2019 IEEE Global Communications Conference (GLOBECOM)*. IEEE, 2019, pp. 1–6.

[11] M. Samir, C. Assi, S. Sharafeddine, D. Ebrahimi, and A. Ghrayeb, "Age of information aware trajectory planning of UAVs in intelligent transportation systems: A deep learning approach," *IEEE Transactions on Vehicular Technology*, vol. 69, no. 11, pp. 12382–12395, 2020.

[12] E. Eldeeb, J. M. de Souza Sant'Ana, D. E. Pérez, M. Shehab, N. H. Mahmood, and H. Alves, "Multi-UAV path learning for age and power optimization in IoT with UAV battery recharge," *IEEE Transactions on Vehicular Technology*, vol. 72, no. 4, pp. 5356–5360, 2022.

[13] R. Kidambi, A. Rajeswaran, P. Netrapalli, and T. Joachims, "Morel: Model-based offline reinforcement learning," in *Advances in Neural Information Processing Systems*, vol. 33. Curran Associates, Inc., 2020, pp. 21810–21823. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2020/file/f7efa4f864ae9b88d43527f4b14f750

[14] K. Ciosek, "Imitation learning by reinforcement learning," in *International Conference on Learning Representations*, 2022. [Online]. Available: https://openreview.net/forum?id=1zwleytEpYx

[15] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, "Trust region policy optimization," in *International conference on machine learning*. PMLR, 2015, pp. 1889–1897.

[16] T. Yu, G. Thomas, L. Yu, S. Ermon, J. Y. Zou, S. Levine, C. Finn, and T. Ma, "Mopo: Model-based offline policy optimization," *Advances in Neural Information Processing Systems*, vol. 33, pp. 14129–14142, 2020.

[17] P. Barde, J. Foerster, D. Nowrouzezahrai, and A. Zhang, "A model-based solution to the offline multi-agent reinforcement learning coordination problem," in *the 23rd International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2024.

[18] K. Yang, C. Shi, C. Shen, J. Yang, S.-p. Yeh, and J. J. Sydir, "Offline reinforcement learning for wireless network optimization with mixture datasets," *IEEE Transactions on Wireless Communications*, pp. 1–1, 2024.

[19] M. G. Bellemare, W. Dabney, and R. Munos, "A distributional perspective on reinforcement learning," in *International conference on machine learning*. PMLR, 2017, pp. 449–458.

[20] W. Dabney, M. Rowland, M. Bellemare, and R. Munos, "Distributional reinforcement learning with quantile regression," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018.

[21] W. Dabney, G. Ostrovski, D. Silver, and R. Munos, "Implicit quantile networks for distributional reinforcement learning," in *International conference on machine learning*. PMLR, 2018, pp. 1096–1105.

[22] R. T. Rockafellar, S. Uryasev *et al.*, "Optimization of conditional value-at-risk," *Journal of risk*, vol. 2, pp. 21–42, 2000.

[23] S. H. Lim and I. Malik, "Distributional reinforcement learning for risk-sensitive policies," *Advances in Neural Information Processing Systems*, vol. 35, pp. 30977–30989, 2022.

[24] Y. Ma, D. Jayaraman, and O. Bastani, "Conservative offline distributional reinforcement learning," in *Advances in Neural Information Processing Systems*, vol. 34. Curran Associates, Inc., 2021, pp. 19235–19247. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2021/file/a05d886123a54de3ca4b0985b718fb9b

[25] J. Jiang and Z. Lu, "Offline decentralized multi-agent reinforcement learning." in *ECAI*, 2023, pp. 1148–1155.

[26] L. Pan, L. Huang, T. Ma, and H. Xu, "Plan better amid conservatism: Offline multi-agent reinforcement learning with actor rectification," in

International Conference on Machine Learning. PMLR, 2022, pp. 17 221–17 237.

[27] J. Shao, Y. Qu, C. Chen, H. Zhang, and X. Ji, "Counterfactual conservative q learning for offline multi-agent reinforcement learning," in *Advances in Neural Information Processing Systems*, vol. 36. Curran Associates, Inc., 2023, pp. 77 290–77 312. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2023/file/f3f2ff9579ba6deeb89caa2fe31b09ff-Paper-Conference.pdf

[28] X. Wang, H. Xu, Y. Zheng, and X. Zhan, "Offline multi-agent reinforcement learning with implicit global-to-local value regularization," in *Advances in Neural Information Processing Systems*, vol. 36. Curran Associates, Inc., 2023, pp. 52 413–52 429. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2023/file/a46c84276e3a4249ab7dbf3e069baf7f-Paper-Conference.pdf

[29] Y. Yang, X. Ma, C. Li, Z. Zheng, Q. Zhang, G. Huang, J. Yang, and Q. Zhao, "Believe what you see: Implicit constraint approach for offline multi-agent reinforcement learning," in *Advances in Neural Information Processing Systems*, vol. 34. Curran Associates, Inc., 2021, pp. 10 299–10 312. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2021/file/550a141f12de6341fba65b0ad0433500-Paper.pdf

[30] E. Eldeeb, M. Shehab, and H. Alves, "Traffic learning and proactive UAV trajectory planning for data uplink in markovian IoT models," *IEEE Internet of Things Journal*, vol. 11, no. 8, pp. 13 496–13 508, 2024.

[31] F. Wu, H. Zhang, J. Wu, L. Song, Z. Han, and H. V. Poor, "AoI minimization for UAV-to-device underlay communication by multi-agent deep reinforcement learning," in *GLOBECOM 2020 - 2020 IEEE Global Communications Conference*, 2020, pp. 1–6.

[32] J. Cui, Y. Liu, and A. Nallanathan, "Multi-agent reinforcement learning-based resource allocation for UAV networks," *IEEE Transactions on Wireless Communications*, vol. 19, no. 2, pp. 729–743, 2020.

[33] Y. S. Nasir and D. Guo, "Multi-agent deep reinforcement learning for dynamic power allocation in wireless networks," *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 10, pp. 2239–2250, 2019.

[34] N. Naderializadeh, J. J. Sydir, M. Simsek, and H. Nikopour, "Resource management in wireless networks via multi-agent deep reinforcement learning," *IEEE Transactions on Wireless Communications*, vol. 20, no. 6, pp. 3507–3523, 2021.

[35] C. Xu, Z. Tang, H. Yu, P. Zeng, and L. Kong, "Digital twin-driven collaborative scheduling for heterogeneous task and edge-end resource via multi-agent deep reinforcement learning," *IEEE Journal on Selected Areas in Communications*, vol. 41, no. 10, pp. 3056–3069, 2023.

[36] Q. Zhang, W. Saad, and M. Bennis, "Millimeter wave communications with an intelligent reflector: Performance optimization and distributional reinforcement learning," *IEEE Transactions on Wireless Communications*, vol. 21, no. 3, pp. 1836–1850, 2021.

[37] ——, "Distributional reinforcement learning for mmwave communications with intelligent reflectors on a uav," in *GLOBECOM 2020-2020 IEEE Global Communications Conference*. IEEE, 2020, pp. 1–6.

[38] Y. Hua, R. Li, Z. Zhao, X. Chen, and H. Zhang, "Gan-powered deep distributional reinforcement learning for resource management in network slicing," *IEEE Journal on Selected Areas in Communications*, vol. 38, no. 2, pp. 334–349, 2019.

[39] P. Sunehag, G. Lever, A. Gruslys, W. M. Czarnecki, V. Zambaldi, M. Jaderberg, M. Lanctot, N. Sonnerat, J. Z. Leibo, K. Tuyls *et al.*, "Value-decomposition networks for cooperative multi-agent learning," *arXiv preprint arXiv:1706.05296*, 2017.

[40] X. Lyu, A. Baisero, Y. Xiao, B. Daley, and C. Amato, "On centralized critics in multi-agent reinforcement learning," 2024. [Online]. Available: https://arxiv.org/abs/2408.14597

[41] R. Lowe, Y. WU, A. Tamar, J. Harb, O. Pieter Abbeel, and I. Mordatch, "Multi-agent actor-critic for mixed cooperative-competitive environments," in *Advances in Neural Information Processing Systems*, vol. 30. Curran Associates, Inc., 2017. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2017/file/68a9750337a418a86fe06c1991a1d64c-Paper.pdf

[42] R. Bellman, "Dynamic programming," *Science*, vol. 153, no. 3731, pp. 34–37, 1966.

[43] H. Xu, L. Jiang, J. Li, Z. Yang, Z. Wang, V. W. K. Chan, and X. Zhan, "Offline rl with no ood actions: In-sample learning via implicit value regularization," *arXiv preprint arXiv:2303.15810*, 2023.

[44] E. Eldeeb, D. E. Pérez, J. Michel de Souza Sant'Ana, M. Shehab, N. H. Mahmood, H. Alves, and M. Latva-Aho, "A learning-based trajectory planning of multiple UAVs for AoI minimization in IoT networks," in *2022 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit)*, 2022, pp. 172–177.

[45] E. Eldeeb, M. Shehab, A. E. Kalø r, P. Popovski, and H. Alves, "Traffic prediction and fast uplink for hidden markov IoT models," *IEEE Internet of Things Journal*, vol. 9, no. 18, pp. 17 172–17 184, 2022.

[46] A. Kosta, N. Pappas, and V. Angelakis, "Age of information: A new concept, metric, and tool," *Foundations and Trends in Networking, Now Publishers, Inc.*, 2017.

[47] E. Eldeeb, M. Shehab, and H. Alves, "Age minimization in massive IoT via UAV swarm: A multi-agent reinforcement learning approach," in *2023 IEEE 34th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, 2023, pp. 1–6.

[48] A. Tampuu, T. Matiisen, D. Kodelja, I. Kuzovkin, K. Korjus, J. Aru, J. Aru, and R. Vicente, "Multiagent cooperation and competition with deep reinforcement learning," *PloS one*, vol. 12, no. 4, p. e0172395, 2017.