

This is an ACCEPTED VERSION of the following published document:

D. Pereira-Ruisánchez, Ó. Fresnedo, D. Pérez-Adán and L. Castedo, "Deep Contextual Bandit and Reinforcement Learning for IRS-Assisted MU-MIMO Systems," in *IEEE Transactions on Vehicular Technology*, vol. 72, no. 7, pp. 9099-9114, July 2023, doi: 10.1109/TVT.2023.3249353.

Link to published version: <https://doi.org/10.1109/TVT.2023.3249353>.

General rights:

© 2023 IEEE. This version of the article has been accepted for publication, after peer review. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The Version of Record is available online at: <https://doi.org/10.1109/TVT.2023.3249353>.

Deep Contextual Bandit and Reinforcement Learning for IRS-assisted MU-MIMO Systems

Dariel Pereira-Ruisánchez, *Student Member, IEEE*, Óscar Fresnedo, *Member, IEEE*,
 Darian Pérez-Adán, and Luis Castedo, *Senior Member, IEEE*

Abstract—The combination of multiple-input multiple-output (MIMO) systems and intelligent reflecting surfaces (IRSs) is foreseen as a critical enabler of beyond 5G (B5G) and 6G. In this work, two different approaches are considered for the joint optimization of the IRS phase-shift matrix and MIMO precoders of an IRS-assisted multi-stream (MS) multi-user MIMO (MU-MIMO) system. Both approaches aim to maximize the system sum-rate for every channel realization. The first proposed solution is a novel contextual bandit (CB) framework with continuous state and action spaces called deep contextual bandit-oriented deep deterministic policy gradient (DCB-DDPG). The second is an innovative deep reinforcement learning (DRL) formulation where the states, actions, and rewards are selected such that the Markov decision process (MDP) property of reinforcement learning (RL) is appropriately met. Both proposals perform remarkably better than state-of-the-art heuristic methods in scenarios with high multi-user interference.

Index Terms—Deep contextual bandit, DDPG, deep reinforcement learning, intelligent reflecting surfaces, MIMO.

I. INTRODUCTION

THE research interest in intelligent reflecting surface (IRS)-assisted multiple-input multiple-output (MIMO) communication systems is enormous nowadays. IRSs are two-dimensional surfaces composed of large arrays of passive scattering elements with specially designed physical structures [1]. Each scattering element can be individually controlled in a software-defined manner to change its electromagnetic (EM) properties and, this way, the phases of the impinging signals. By properly designing the MIMO precoding/combining matrices and the IRS phase-shift matrix, the demanding requirements of the use cases in beyond 5G (B5G) and sixth generation (6G) systems are expected to be accomplished [2]–[7].

In conventional MIMO systems, the communication channel is an uncontrollable factor with a high impact on the reception quality. In this sense, IRS-assisted communications are the most appealing alternative to configure the propagation environment smartly [2]. In [8], the authors analyze the deployment of IRS-assisted multi-user MIMO (MU-MIMO) systems as an energy-efficient alternative to multi-antenna amplify-and-forward relaying. Numerical results show that the IRS-based approach provides up to 300% higher energy efficiency (EE) while reaching near optimal spectrum efficiency (SE). Authors in [9] address the SE performance of an IRS-assisted multi-user multiple-input single-output (MU-MISO) downlink communication in a scenario with non-line-of-sight (NLoS) conditions between the base station (BS) and the users. As

expected, the sum-rate values significantly increase when considering the BS-IRS-user channel paths.

Due to the dynamic nature of vehicular communications, IRS is a promising technology for achieving cost and energy-efficient communications via smartly reshaping the wireless propagation channel [10]. According to [11], unmanned aerial vehicle (UAV) applications and cellular vehicular communications enhance by deploying IRSs. For instance, IRSs have been considered both to improve the reliability of the ground-to-air links and to be UAV-carried to provide controllable on-demand coverage [12].

Many challenges related to the deployment of IRS-assisted MIMO systems remain unsolved [13], [14]. Control and deployment complexities are two major issues that increase with the number of configurable elements in both technologies. In IRS-assisted systems, the average signal-to-noise ratio (SNR) and the transmission power savings are directly related to the square of the number of scattering elements [2]. In MIMO systems, the multiplexing gain, the spatial diversity, and the favorable propagation conditions improve when considering more antennas [15]. However, in MIMO and IRS-assisted communications, increasing the number of configurable elements leads to optimization problems that are often too computationally expensive and incur considerable timing overhead.

On the one hand, the complexity of the optimization problems when considering IRS-assisted communications increases because of some characteristics of the sets of feasible solutions. The passive behavior of the scattering elements in IRSs is usually modeled through complex-valued entries with modulus constraints, which are difficult to handle and lead to non-convex formulations of the problem. On the other hand, the requirement of jointly optimizing the users' precoders and the IRS phase-shifts in the IRS-assisted MIMO systems leads to optimization problems with a non-convex cost function over a non-convex search space. Some conventional model-based approaches relax the original problem to obtain a convex formulation, approximate the non-convex constraints iteratively, develop heuristic solutions or change the optimization domain [4], [16]–[22]. However, in those cases, optimality is sacrificed for tractability. In addition, due to the increase in complexity and heterogeneity of these communications, the solutions to the optimization problems cannot wholly rely on models, which are either too computationally demanding or inconsistent with the actual behavior of practical deployments. Among others, these issues render conventional model-based approaches inefficient for most emerging time-sensitive applications.

Several other approaches consider data-driven alternatives to deal with the high complexity of designing IRS-assisted MIMO systems. Supervised deep learning (DL) has been widely considered to solve problems whose conventional solutions become too computationally complex, or the search spaces are too vast for considering approaches like genetic algorithms [23]–[27]. However, the performance of supervised DL is highly dependent on the data sets used for training, and the predictions are very sensitive to modifications in the communication systems [28].

Due to the limitations of the solutions above, recent works suggest the use of deep reinforcement learning (DRL). Unlike supervised DL, DRL performs the training by interacting with the communication system. Hence, DRL approaches do not require massive labeled sets of data. DRL-based solutions have been employed in different scenarios for wireless communication systems, proving to be an appealing alternative when tackling complex optimization problems [29]–[33]. However, defining a Markov decision process (MDP)—a requirement for reinforcement learning (RL) formulations—makes no sense for some optimization problems, and simpler formulations such as contextual bandit (CB) could be more appropriate.

According to [34], CB problems are intermediate between the k -armed bandit and RL problems. CB involves learning optimal actions but aiming to maximize an immediate reward instead of a long-term one. A significant difference with RL is that actions and states in CB are formulated so that actions do not affect the following states or rewards.

Recently, CB has been gaining attention, and several CB-based solutions have been proposed for different wireless communication open problems. In [35], the authors propose a deep CB (DCB)-based instantaneous beam selection method for mm-wave (mmWave) cell-free networks. The results show that the proposed method outperforms other state-of-the-art algorithms in terms of latency and computational complexity, thus enabling faster initial access to the network. Authors in [36] proposed an algorithm based on CB for beam width selection in mmWave links. Their proposal efficiently handled the fast channel variations by considering a codebook-aided approach. Authors in [37] also propose a CB-based solution for mmWave device-to-device two-hop relay probing, which outperforms the considered benchmarks in terms of relaying performance and execution time.

CB-based algorithms are an appealing alternative to handle the challenging scenarios of B5G and 6G. However, as explained in [38], standard CB settings describe problems where optimal actions are selected from discrete sets of feasible actions. Hence, most existing works consider only discrete action spaces, which are unsuitable for joint optimization of the IRS phase-shift matrix and the MIMO precoders. This optimization problem becomes a combinatorial search problem with an intractable number of possible actions when considering a discrete set of actions. Additionally, unlike standard CB settings, the actions in this optimization problem are expected to be composed of several configurable elements that must be jointly optimized. Therefore, the action space in this problem is high-dimensional, in the sense that many individual elements must be optimized, and continuous because each

element is continuous-valued. These issues and the limitations of the scarce solutions in the literature have motivated our work, whose main contributions are:

- The proposal of a new approach for the joint optimization of the IRS phase-shift and MIMO precoder matrices, considering a CB formulation of the problem with continuous action and state spaces.
- The development of an actor-critic framework called deep contextual bandit-oriented deep deterministic policy gradient (DCB-DDPG), which enables handling the high-dimensional continuous action space by considering a deterministic policy gradient approach.
- The use of a multi-head structure for the actor artificial neural network (ANN) in the DCB-DDPG framework. This way, the precoders' normalization, and the IRS phase-shift matrix projection stages are performed within the actor. This proposal remarkably improves the training performance and is a major step for multi-agent implementations.
- The design of a DRL-based deep deterministic policy gradient (DDPG) formulation of the joint optimization problem where state and action spaces are selected in a way that is more appropriate to fit the requirements of the MDP.
- The evaluation of the performance of the proposed methods in an IRS-assisted multi-stream (MS) MU-MIMO system. We particularly focus on strong interference communication scenarios where the number of transmitted streams is larger than the number of receiving antennas, and the receiver does not have enough degrees of freedom to manage the multi-user interference.

Hence, we propose an innovative approach to a well-known DRL-based formulation (DRL-DDPG) and a game-changing approach based on DCB (DCB-DDPG). With a detailed analysis of both frameworks, we intend to understand better their strengths and weaknesses beyond those that can be inferred from observing the simulation results. The remainder of this paper is structured as follows. Section II details the theoretical fundamentals of the CB and RL frameworks and briefly analyzes some related works. Section III introduces the IRS-assisted MS MU-MIMO system model and formulates the optimization problem to determine the IRS phase-shift matrix and MIMO precoders. Sections IV and V explain the CB and RL proposed solutions, respectively. In Section VI, we present an analysis of their convergence and computational complexity. Section VII is devoted to the simulation results, and Section VIII presents the conclusions.

II. PRELIMINARIES

A. Reinforcement Learning (RL) and Contextual Bandit (CB)

As stated in [34], RL is a computational approach to learning-by-interacting, i.e., mapping situations to actions that maximize a numerical reward function. Most RL problems are formalized in terms of MDPs, where the agent (the learning and decision-maker element) interacts with the environment (all the elements external to the agent) through actions that affect the following states and rewards. Hence, RL problems




Fig. 1. Contextual bandit (CB) framework.

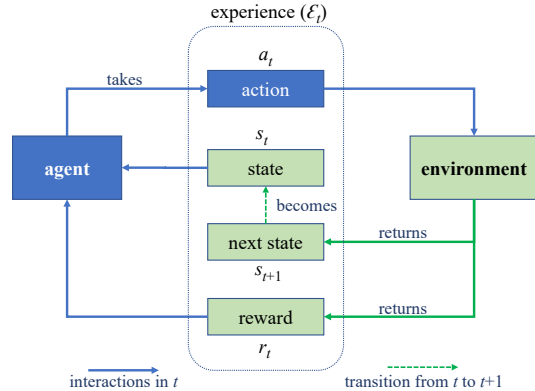


Fig. 2. Reinforcement learning (RL) framework.

are described by a dynamics function $p(s_{t+1}, r_t | a_t, s_t)$ such that in every time instant t , the next state s_{t+1} , and the reward r_t are conditioned by the effect of taking an action a_t in the current state s_t .

On the other hand, CB problems can be understood as a relaxation of RL problems. In CB, interactions fit a dynamics function $p(r_t | a_t, s_t)$. Thus, only immediate rewards are affected by the current state and action. Unlike RL, a state in CB only contains the necessary information to select the best action for that state.

Fig. 1 and Fig. 2 show the interactions between the main components of the CB and RL frameworks, respectively. The environment is the external element with which the agents interact and sets the state s_t in every instant t . However, in the CB formulation, the selection of the current state is not related to the previous state and action. In both approaches, the agent selects the action a_t that maximizes the value of a reward function, i.e., the immediate reward $r_t(\cdot)$ in CB and a function of the expected long-term reward $q_\pi(\cdot)$ in RL. The tuples $\mathcal{E}_t = (s_t, a_t, r_t)$ in CB and $\mathcal{E}_t = (s_t, a_t, r_t, s_{t+1})$ in RL are commonly termed experiences.

The policy is another critical element of both frameworks. It is the decision-making rule that defines the action that the agent will take while being in a given state, i.e., $a_t = \pi(s_t)$, where a_t belongs to the set $\mathcal{A}(s_t)$ of available actions in s_t , and $\pi(\cdot)$ is a deterministic policy. Hence, CB and RL algorithms aim to find the optimal policy $\pi^*(\cdot)$ that maximizes the respective reward functions.

Conventional tabular approaches to CB and RL problems have proven efficient when considering low-dimensional, dis-

crete state and action spaces [34]. However, these algorithms rapidly become intractable in problems where action and state spaces are continuous or arbitrarily large. In this regard, the ANN-based DCB and DRL algorithms are appealing alternatives. The use of ANNs for function approximation of the policy and reward functions enables a wide range of new approaches to high-complexity problems like the one we are addressing in this work.

B. Related Works

In this subsection, we analyze two existing approaches to the joint optimization of the IRS and precoder matrices. We focus on [39] and [40], which formulate the joint optimization as an RL problem and a CB problem, respectively. Table I summarizes some differences and similarities between these two approaches.

Authors in [39] propose an RL-based DDPG framework for the joint optimization of the IRS matrix and the precoder in the downlink of a single-stream MU-MISO system. In their proposal, the action vectors are composed of the entries of the IRS and precoder matrices. These matrices are considered continuous-valued, enabling high flexibility in searching for the optimal solution. Besides, it allows authors to overcome the high dimensionality of discrete-valued implementations. The instantaneous sum-rate value is the reward function, matching the metric considered for the optimization problem.

As shown in Table I, the received and transmitted powers are included in the states [39]. However, states also include the previous entries of the IRS phase-shift and precoder matrices. Since the reward is a function of the current sum-rate value, it is unclear how the selection of the optimal matrices for a given channel realization could be related to the previously employed matrices. In RL, the agent takes the action according to the state information. Hence, for a given channel realization, there would be as many possible optimal matrices as possible actions in the previous state. It seems more accurate to consider actions that only depend on the current channel realization and therefore rewards that only depend on the current channel realization and the IRS and precoding matrices.

On the other hand, [40] considers a CB-based approach to the joint optimization of the precoders and the IRS matrices in the downlink of a multi-IRS MU-MISO system. As in [39], the instantaneous sum-rate value is the reward. However, in [40], the state vectors are only composed of the channel coefficients for the current channel realization, which properly follows the CB formulation.

The major drawback of the solution proposed in [40] is the assumption of only discrete-valued actions, i.e., IRS scattering elements with only two possible phase values and a codebook set for the precoder selection. In setups like that, with low-dimensional sets of possible actions, the system performance is limited, and the benefits of IRS-assisted MIMO communications cannot be fully leveraged. On the other hand, the complexity of the proposed algorithm steeply increases with the number of possible actions. This issue is interpreted in [34], [41] as an instance of the curse of dimensionality.

TABLE I
RELATED WORKS' SUMMARY

Reference	[39]	[40]
Problem formulation	Reinforcement learning	Contextual bandit
State space	Previous time step IRS matrix values, previous time step precoder values, CSI and power-related values (Continuous-valued)	CSI (Continuous-valued)
Action space	Updated IRS matrix and precoder (Continuous-valued)	Updated IRS matrices and precoder (Discrete-valued)

C. Overcoming the Limitations

Based on the previous analysis of related works, we propose two different approaches to overcome their limitations. The first is a CB formulation of the joint optimization problem where action and state spaces are considered to be continuous. To handle this situation, we introduce a novel framework called DCB-DDPG. It leverages several features of DDPG and adapts them to the CB formulation. As introduced in [41], DDPG is a model-free, off-policy actor-critic framework based on the deterministic policy gradient (DPG) algorithm. DDPG has rapidly become a well-established solution to high-dimensional, continuous action spaces and its features make it a suitable alternative to assist our DCB formulation of the joint optimization problem [39], [42]–[45].

In this sense, model-free algorithms enable learning from interactions in scenarios where system models become too complex or distant from the actual behavior of practical deployments. Although the initial training of the proposed solution is based on the outcomes of the IRS-assisted MS MU-MIMO uplink system model, its performance during practical evaluations will depend on its capability to keep learning from interacting. Besides, the off-policy behavior enables exploring more efficiently high-dimensional search spaces like ours and using past experiences during training. Off-policy algorithms can be trained with either externally generated or self-generated experiences, thus reducing the possibility of getting stuck in a local minimum [34]. Finally, the DPG updates are suitable for our aim of finding a deterministic policy that returns the continuous-valued optimal action in every continuous-valued state while keeping the evaluation variance as low as possible [46].

From the above analysis, the second proposal is a DRL-based DDPG approach. However, we consider sets of states and actions that differ from those considered in the related work [39] and, as discussed in later sections, can be more appropriate to fit the requirements of the MDP. For greater clarity, this second framework will be termed DRL-DDPG.

During the initial phases of the investigation, we analyzed other DRL-based algorithms that lacked some of the desired features above, i.e., they were either on-policy methods or learned stochastic policies (e.g., proximal policy optimization (PPO), trust region policy optimization (TRPO) and asynchronous advantage actor-critic (A3C) algorithms). In on-policy methods, exploration is restricted to the learned target policy plus a certain degree of randomness in the actions taken, which is a limiting factor in high-dimensional continuous-valued search spaces [34]. On the other hand, stochastic policies are more desirable in partially observable

MDP (POMDP) formulations of RL problems, where the environment uncertainty characterizes the interactions [34], [47]. However, in our formulation, the relationship between the observed states, actions, and rewards is established deterministically: system sum-rate values are unequivocally related to the channel realizations and the configured IRS matrix and precoders. In addition, to improve the performance of our deterministic approach, we consider several mechanisms that enable efficient exploration of both the action and the state spaces.

For a practical evaluation of both proposed approaches, the considered system model has been extended to MS MU-MIMO communications. Additionally, we address several scenarios where there are not enough degrees of freedom to cancel all the multi-user interference.

III. SYSTEM MODEL AND OPTIMIZATION PROBLEM

A. Notation

Along this work, the following notation will be employed: a is a scalar, \mathbf{a} is a vector, and \mathbf{A} represents a matrix. Transpose, conjugate transpose, and the Frobenius norm of \mathbf{A} are represented by \mathbf{A}^T , \mathbf{A}^H , and $\|\mathbf{A}\|_F^2$, respectively. Calligraphic letters are employed to denote sets and tuples. $|\mathcal{R}|$ stands for the cardinality of a set \mathcal{R} . \mathbf{I}_N indicates an $N \times N$ identity matrix, and \mathcal{I}_N denotes the set of integers from 1 to N . The operator $\text{blkdiag}(\cdot)$ constructs a block diagonal matrix from its input matrices. Finally, $\text{flatten}(\cdot)$ is the operator that reshapes any matrix $\mathbf{V} \in \mathbb{C}^{A \times B}$ into a vector $\mathbf{v} \in \mathbb{C}^{1 \times AB}$ by concatenating all the entries. The mathematical relationships presented in the following sections hold for all the consecutive time steps t that fit within one coherence block. Hence, for the sake of simplicity, sub-index t is used only where necessary to avoid ambiguities.

B. IRS-assisted MS MU-MIMO Uplink

Let us consider the uplink of an IRS-assisted MS MU-MIMO communication system, as shown in Fig. 3. In this scenario, each of the K users employs N_t antennas to send N_s data streams to a BS equipped with N_r antennas. We assume there is a total blockage between the users and the BS, i.e., no direct channel is available between any of the K users and the BS.

The symbols transmitted by the k -th user on each channel usage are represented by $\mathbf{x}_k = [x_{k,1}, \dots, x_{k,N_s}]^T \in \mathbb{C}^{N_s \times 1}$. We assume \mathbf{x}_k follows a zero-mean multivariate complex-valued Gaussian distribution, i.e., $\mathbf{x}_k \sim \mathcal{N}_{\mathbb{C}}(0, \mathbf{I}_{N_s})$. Such symbols are linearly processed with the precoder $\mathbf{P}_k \in \mathbb{C}^{N_t \times N_s}$.




Fig. 3. Uplink of an IRS-assisted MS MU-MIMO system.

The power transmitted by the k -th user is limited to Ω_k , which leads to the individual power constraint on the precoders given by $\|\mathbf{P}_k\|_F^2 \leq \Omega_k, \forall k$.

The deployed IRS is assumed to have N scattering elements. Thus, the IRS phase-shift matrix is represented by the diagonal matrix $\Theta = \text{diag}(e^{j\theta_1}, \dots, e^{j\theta_N}) \in \mathcal{D}$ where $\theta_n \in [0, 2\pi)$ represents the phase shift introduced by the n -th element of the IRS. $\mathcal{D} \in \mathbb{C}^{N \times N}$ is the set of feasible IRS matrices, i.e., the diagonal matrices with unit modulus entries. The vector that contains the entries in the main diagonal of Θ is denoted by $\theta \in \mathbb{C}^{N \times 1}$.

According to this system model, the received signal at the BS is given by

$$\mathbf{y} = \mathbf{H}_{\text{IB}} \Theta \sum_{k=1}^K \mathbf{H}_{\text{UI}k} \mathbf{P}_k \mathbf{x}_k + \mathbf{n}, \quad (1)$$

where the matrix $\mathbf{H}_{\text{UI}k} \in \mathbb{C}^{N \times N_t}$ represents the channel response of the link from the k -th user to the IRS, $\mathbf{H}_{\text{IB}} \in \mathbb{C}^{N_r \times N}$ stands for the channel response from the IRS to the BS, and $\mathbf{n} = [n_1, n_2, \dots, n_{N_r}]^T$ represents the complex-valued additive white Gaussian noise (AWGN) which is modeled as $\mathbf{n} \sim \mathcal{N}_{\mathbb{C}}(0, \sigma_n^2 \mathbf{I}_{N_r})$.

Using a more compact notation, the received signal at the BS given by (1) can be rewritten as

$$\mathbf{y} = \mathbf{H}_{\text{IB}} \Theta \mathbf{H}_{\text{UI}} \mathbf{P} \mathbf{x} + \mathbf{n}, \quad (2)$$

where $\mathbf{H}_{\text{UI}} = [\mathbf{H}_{\text{UI}1}, \dots, \mathbf{H}_{\text{UI}K}] \in \mathbb{C}^{N \times KN_t}$, $\mathbf{x} = [\mathbf{x}_1^T, \dots, \mathbf{x}_K^T]^T \in \mathbb{C}^{KN_s \times 1}$ and $\mathbf{P} = \text{blkdiag}(\mathbf{P}_1, \dots, \mathbf{P}_K) \in \mathcal{P}$ is the block-diagonal matrix which stacks all the users' precoders. $\mathcal{P} \in \mathbb{C}^{N_t K \times N_s K}$ is the set of block-diagonal matrices with norm-constrained submatrices.

The vector with all the estimated user symbols $\hat{\mathbf{x}} = [\hat{\mathbf{x}}_1^T, \dots, \hat{\mathbf{x}}_K^T]^T \in \mathbb{C}^{KN_s \times 1}$ is obtained by the linear filtering of the signal received at the BS, i.e., $\hat{\mathbf{x}} = \mathbf{W}^H \mathbf{y}$, where $\mathbf{W}^H = [\mathbf{W}_1^H, \dots, \mathbf{W}_K^H]^H \in \mathbb{C}^{KN_s \times N_r}$ is the BS receiving filter matrix that stacks the individual receiving filter matrices $\mathbf{W}_k^H \in \mathbb{C}^{N_s \times N_r}$. We assume $\mathbf{W}_k^H, \forall k$ to be the minimum mean square error (MMSE) receiving filter given by

$$\mathbf{W}_k^H = \mathbf{P}_k^H \mathbf{H}_{\text{UI}k}^H \Theta^H \mathbf{H}_{\text{IB}}^H \times (\mathbf{H}_{\text{IB}} \Theta \mathbf{H}_{\text{UI}} \mathbf{P} \mathbf{P}^H \mathbf{H}_{\text{UI}}^H \Theta^H \mathbf{H}_{\text{IB}}^H + \sigma_n^2 \mathbf{I}_{N_r})^{-1}. \quad (3)$$

 TABLE II
SYSTEM MODEL PARAMETERS

Parameter	Description
K	number of users
N_t	number of transmitting antennas per user
N_s	number of transmitted streams per user
N_r	number of receiving antennas at BS
\mathbf{H}_{UI}	channel response matrix from the users to the IRS
\mathbf{H}_{IB}	channel response matrix from the IRS to the BS
\mathbf{P}_k	precoder matrix for each user k
\mathbf{P}	block-diagonal matrix with all the users' precoders
Θ	IRS phase-shift matrix
θ	vector of the IRS phase-shifts
Ω_k	users' individual power constraint
\mathbf{W}_k^H	MMSE receiving filter for each user k
\mathbf{X}_k	interference plus noise matrix for each user k
R_k	achievable rate for user k

Considering this system model, we look for the IRS phase-shift matrix and the precoders that maximize the achievable sum-rate. Towards this aim, we formulate the following optimization problem

$$\begin{aligned} & \arg \max_{\mathbf{P}, \Theta} \sum_{k=1}^K R_k \\ & \text{s.t. } \|\mathbf{P}_k\|_F^2 \leq \Omega_k, \forall k \in \mathcal{I}_K, \\ & \quad \Theta \in \mathcal{D}, \end{aligned} \quad (4)$$

where R_k is the k -th user individual rate which is given by

$$R_k = \log_2 \det (\mathbf{I}_K + \mathbf{X}_k^{-1} \times \mathbf{W}_k^H \mathbf{H}_{\text{IB}} \Theta \mathbf{H}_{\text{UI}k} \mathbf{P}_k \mathbf{P}_k^H \mathbf{H}_{\text{UI}k}^H \Theta^H \mathbf{H}_{\text{IB}}^H \mathbf{W}_k), \quad (5)$$

where

$$\begin{aligned} \mathbf{X}_k &= \sum_{i \neq k} \mathbf{W}_k^H \mathbf{H}_{\text{IB}} \Theta \mathbf{H}_{\text{UI}i} \mathbf{P}_i \mathbf{P}_i^H \mathbf{H}_{\text{UI}i}^H \Theta^H \mathbf{H}_{\text{IB}}^H \mathbf{W}_k \\ & + \sigma_n^2 \mathbf{W}_k^H \mathbf{W}_k \end{aligned} \quad (6)$$

is the interference plus noise matrix, and \mathbf{W}_k^H is the MMSE individual receiving filter for the k -th user, as in (3). Notice that for the single-stream scenarios, $N_s = 1$, the user individual receiving filter matrices $\mathbf{W}_k^H \in \mathbb{C}^{N_s \times N_r}$ collapse in the vectors $\mathbf{w}_k^H \in \mathbb{C}^{1 \times N_r}$.

As stated in (4), the precoders must meet the individual power constraints $\|\mathbf{P}_k\|_F^2 \leq \Omega_k, \forall k$, where Ω_k represents the available power at the k -th user. For the sake of simplicity, we assume the same power constraint value for all the users, i.e., $\Omega_k = \Omega, \forall k$. Without loss of generality, we also assume a noise variance σ_n^2 equal to one. Therefore, the SNR per user is given by $\text{SNR}_{(\text{dB})} = 10 \log_{10}(\Omega)$.

Table II summarizes the main system model parameters and their descriptions. Due to the choice of the objective function and the design constraints for the IRS and the users' precoders, (4) becomes a non-convex, non-trivial optimization problem. We propose using the DCB-DDPG and DRL-DDPG frameworks to handle this problem efficiently. Although we will only address uplink communication scenarios like the one previously described, proposed CB and RL solutions can be easily adapted to downlink scenarios, and similar results are obtained.

C. Channel Model

Channel responses from the users to the IRS ($\mathbf{H}_{\text{UI}} = [\mathbf{H}_{\text{UI}1}, \dots, \mathbf{H}_{\text{UI}K}]$) and from the IRS to the BS (\mathbf{H}_{IB}) are assumed to be perfectly known, as in [30], [31], [35], [40], [48]. Although channel acquisition is challenging in IRS-aided systems, a promising line of research is committed to facing this problem [49]–[51]. Therefore, it is reasonable to assume that the IRS-assisted communication systems can implement some appropriate mechanism to provide accurate channel state information (CSI) to the agents in the DCB and RL frameworks.

In particular, we assume a Rayleigh fading model for $\mathbf{H}_{\text{UI}k}$, i.e., the entries of $\mathbf{H}_{\text{UI}k}$, $\forall k$ are independent and identically distributed (i.i.d.) complex-valued circularly symmetric Gaussian random variables. On the other hand, considering that IRSs are typically installed to ensure line-of-sight (LoS) to the BS, a Rician fading channel model is adopted to describe \mathbf{H}_{IB} as in [40], [48], i.e.,

$$\mathbf{H}_{\text{IB}} = \sqrt{\frac{\beta}{1+\beta}} \mathbf{H}_{\text{IB}}^{\text{LOS}} + \sqrt{\frac{1}{1+\beta}} \mathbf{H}_{\text{IB}}^{\text{NLOS}}, \quad (7)$$

where β is the Rician factor, which is set to 3 dB. $\mathbf{H}_{\text{IB}}^{\text{LOS}}$ and $\mathbf{H}_{\text{IB}}^{\text{NLOS}}$ stand for the LoS component and the Rayleigh fading component, respectively.

IV. CB-BASED JOINT OPTIMIZATION

In this section, the optimization problem in (4) is solved using a novel CB-based approach. We start selecting the state, action, and reward spaces. Next, we review some essential components in the proposed DCB-DDPG framework, and finally, we derive the algorithmic solution.

A. DCB-DDPG: State, Action and Reward

According to the CB formulation, we introduce the following states, actions, and rewards:

- The **state** vector \mathbf{s}_t comprises the current values of all the channel response matrices from the users to the IRS ($\mathbf{H}_{\text{UI}k}$, $\forall k$) and from the IRS to the BS (\mathbf{H}_{IB}). The state vector is constructed such that

$$\mathbf{s}_t = [\text{flatten}(\mathbf{H}_{\text{UI}1}), \dots, \text{flatten}(\mathbf{H}_{\text{UI}K}), \text{flatten}(\mathbf{H}_{\text{IB}})]. \quad (8)$$

Hence, the dimension of the state space vectors is $D_{\text{state}}^{\text{DCB}} = KN_tN + N_rN$. Recall that the state space is continuous since the entries of \mathbf{s}_t can take any complex value. Notice that we assume signal processing techniques that enable handling complex-valued entries. Otherwise, the imaginary and real parts must be treated as independent inputs, leading to vectors twice the size.

- The **action** vector \mathbf{a}_t is composed of the entries in the main diagonal of the IRS phase-shift matrix ($\boldsymbol{\theta}$) and those in all the users' precoders (\mathbf{P}_k , $\forall k$). Hence, the action vector is constructed such that

$$\mathbf{a}_t = [\text{flatten}(\mathbf{P}_1), \dots, \text{flatten}(\mathbf{P}_K), \text{flatten}(\boldsymbol{\theta})]. \quad (9)$$

The dimension of the action space vectors is $D_{\text{action}}^{\text{DCB}} = KN_tN_s + N$. The action space is also continuous, though the matrices formed by these entries must meet the previously defined constraints ($\|\mathbf{P}_k\|_{\text{F}}^2 \leq \Omega$, $\forall k$, and $\boldsymbol{\Theta} \in \mathcal{D}$).

- The **reward** r_t is equal to the system sum-rate determined according to (4), and by considering the entries of the current state and the action vectors, i.e.,

$$r_t = \sum_{k=1}^K R_k(\mathbf{s}_t, \mathbf{a}_t). \quad (10)$$

B. DCB-DDPG: Framework Elements

CB algorithms are mostly oriented to solve discrete action space formulations, such as those presented in [35]–[37], [40]. Hence, to address the high-dimensional continuous action space of the CB-based formulation presented in the previous subsection, we have developed the proposed DCB-DDPG framework. Fig. 4 shows a representation of the agent that we consider in our proposal. Some of the elements that support this setup are the following:

- **Actor-critic**: called this way due to the interactions between the reward and policy functions. The critic is the element that updates the reward function according to the observed instantaneous rewards. Meanwhile, the actor updates the policy function by learning the actions that maximize the reward function. The update of the actor function is made through deterministic policy gradient updates as described in [41], [46]. Iteratively, both converge to the corresponding optimal function.
- **ANNs for function approximation**: the reward and the policy functions are represented through ANNs. This way, the continuous action and state spaces can be efficiently handled. The actor network $\pi(\mathbf{s}, \boldsymbol{\vartheta}_\pi)$ is the approximation of the policy function, and $\boldsymbol{\vartheta}_\pi$ is the vector of weights of the actor network to be trained. Similarly, $r(\mathbf{s}, \mathbf{a}, \boldsymbol{\vartheta}_r)$ is the approximation of the reward function, and $\boldsymbol{\vartheta}_r$ is the vector of weights of the critic network.
- **Experience replay**: the experience tuples \mathcal{E}_t observed at each time step are stored in a replay memory \mathcal{R} , such that they can be used for off-policy training. $|\mathcal{R}|$ is the number of experience tuples stored, and random mini-batches \mathcal{B} with size $|\mathcal{B}|$ are used to update the weights of the actor and critic networks during training. Using the experience replay improves the stability of the learning stages and eliminates the undesirable effects of the correlation between samples.

C. DCB-DDPG: Proposed Algorithm

In this section, we present the algorithm proposed to train the DCB-DDPG agent for solving the optimization problem in (4). The values of the IRS phase-shift matrix ($\boldsymbol{\Theta}$) and the precoder (\mathbf{P}) that maximize the achievable sum-rate are learned by following the steps in Algorithm 1.

We first create the actor and critic networks by randomly initializing the network parameters ($\boldsymbol{\vartheta}_\pi$ and $\boldsymbol{\vartheta}_r$, respectively). We also create the actor network optimizer and the critic

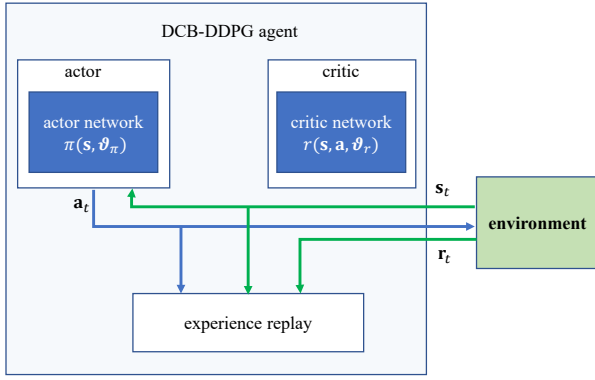


Fig. 4. Elements in the DCB-DDPG agent.

Algorithm 1 DCB-DDPG algorithm

-
- 1: **Initialize:**
 - 2: set $\pi(\mathbf{s}, \boldsymbol{\vartheta}_\pi)$ given random $\boldsymbol{\vartheta}_\pi$
 - 3: set $r(\mathbf{s}, \mathbf{a}, \boldsymbol{\vartheta}_r)$ given random $\boldsymbol{\vartheta}_r$
 - 4: create \mathcal{R}
 - 5: **for** $t = 0, \dots, T - 1$ **do:**
 - 6: set \mathbf{s}_t given channels $\mathbf{H}_{\text{UI}k}, \forall k$ and \mathbf{H}_{IB}
 - 7: agent takes $\mathbf{a}_t = \pi(\mathbf{s}_t, \boldsymbol{\vartheta}_\pi) + \mathbf{n}_e$
 - 8: environment returns r_t
 - 9: \mathcal{R} stores $\mathcal{E}_t = (\mathbf{s}_t, \mathbf{a}_t, r_t)$
 - 10: **if** $|\mathcal{R}| > |\mathcal{B}|$:
 - 11: agent samples $|\mathcal{B}|$ random experiences:
 - 12: $\mathcal{E}_i = (\mathbf{s}_i, \mathbf{a}_i, r_i), i = 0, \dots, |\mathcal{B}| - 1$
 - 13: calculate the critic loss L_c by using (11)
 - 14: update $\boldsymbol{\vartheta}_r$ by the back-propagation of L_c
 - 15: calculate the actor loss L_a by using (12)
 - 16: update $\boldsymbol{\vartheta}_\pi$ by the back-propagation of L_a
 - 17: **end if**
 - 18: Obtain \mathbf{P} and $\boldsymbol{\Theta}$ by evaluating policy $\pi(\mathbf{s}, \boldsymbol{\vartheta}_\pi)$
-
- Output:** trained actor network $\pi(\mathbf{s}, \boldsymbol{\vartheta}_\pi)$
-

network optimizer to handle the updates of $\boldsymbol{\vartheta}_\pi$ and $\boldsymbol{\vartheta}_r$, respectively. The experience replay buffer \mathcal{R} is created with no elements.

We propose to split the interactions between the DCB-DDPG agent and the environment into T time steps. Since no terminal states can be defined in our problem, the number of time steps can be arbitrarily selected. The state \mathbf{s}_t is created as in (8) from randomly generated $\mathbf{H}_{\text{UI}k}, \forall k$ and \mathbf{H}_{IB} channel response matrices. Hence, episodes are not required in this CB formulation since the randomness of the channel response matrices guarantees continuous exploration over the state space.

From the current state \mathbf{s}_t , the agent selects the action \mathbf{a}_t according to the output of the actor network and an exploration noise $\mathbf{n}_e \sim \mathcal{N}_{\mathbb{C}}(0, \sigma_{\mathbf{n}_e}^2 \mathbf{I}_{D_{\text{action}}^{\text{DCB}}})$. Incorporating this noise improves the exploration of the action space since actions around the one selected by the actor network are evaluated. As explained in [34], [41], [52], in off-policy algorithms, the policy resulting from adding the noise can be treated as a different policy. Samples for training can be obtained by following any policy, and the learning capability is not

affected.

The learning stage (lines 11 to 16 in Algorithm 1) is first performed when there are enough stored experiences to sample a mini-batch of size $|\mathcal{B}|$. Notice that the first mini-batches will have several samples in common since $|\mathcal{B}| \approx |\mathcal{R}|$. However, this effect soon disappears as $|\mathcal{R}|$ increases in every time step.

The experience tuples in \mathcal{B} are employed to train the actor and critic networks as described in lines 13 to 16. This algorithm aims to maximize the instantaneous reward, not a long-term reward function, as in RL. Hence, the sampled reward values are used as targets for the critic training. The critic loss, L_c , is then calculated as

$$L_c = \frac{1}{|\mathcal{B}|} \sum_i (r_i - r(\mathbf{s}_i, \mathbf{a}_i, \boldsymbol{\vartheta}_r))^2 \quad (11)$$

and the obtained values are back-propagated by using the critic network optimizer. This way, the critic network is trained to minimize L_c and represent the behavior of the actual reward function as closely as possible.

Later, by following a deterministic policy gradient update, the actor network is trained to predict the action that maximizes the output of the critic network. Because of this, the sign of the critic output value is changed, i.e.,

$$L_a = -\frac{1}{|\mathcal{B}|} \sum_i r(\mathbf{s}_i, \pi(\mathbf{s}_i, \boldsymbol{\vartheta}_\pi), \boldsymbol{\vartheta}_r), \quad (12)$$

and, hence, the back-propagation of the actor network optimizer can be perceived as a gradient ascent instead of a descent. Notice that L_a is intended only to update the actor network parameters. Thus, the actor network optimizer does not update the critic network parameters.

At the end of every learning time step, the performance of the obtained policy is tested by running an evaluation step. Only unseen channel realizations are visited during the evaluation, and no learning or exploration is performed. The evaluation stage has two main aims: the detection of undesired behaviors during the learning stage, which could lead to performance degradation, and the identification of the policies that perform better with the unseen channel responses so that they can be used in the operational stage. The expected result is to obtain a trained policy network capable of predicting near-optimal IRS and precoder matrices.

D. DCB-DDPG: ANNs Structure

Fig. 5 shows the structures of the ANNs that we propose to use as the actor and critic networks. The actor network input and output layer dimensions equal $D_{\text{state}}^{\text{DCB}}$ and $D_{\text{action}}^{\text{DCB}}$, respectively. For this network, we employ a multi-head ANN setup such that each user precoder's normalization and the projection of the phase-shift values into unit modulus values are handled in independent normalization/projection layers. This structure enables high flexibility in power allocation since individual constraints can be assigned to the users. Besides, it is an appealing choice for multi-agent implementations since the backbone network can be trained centrally and shared with all the users and the IRS controller. The hidden layers are fully connected layers. They are composed of $2D_{\text{state}}^{\text{DCB}}$

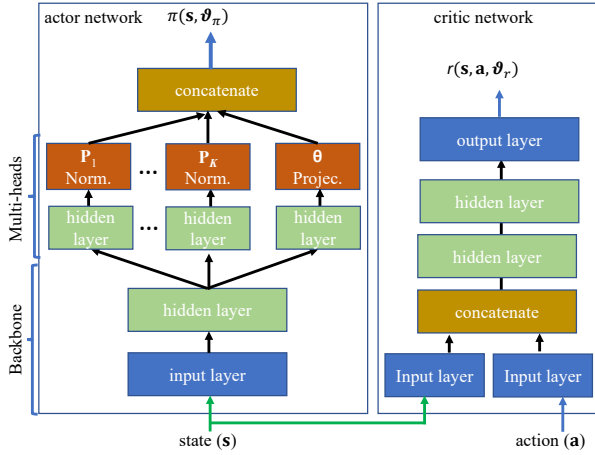


Fig. 5. DCB-DDPG actor and critic network structure.

output neurons in the network backbone, $N_t N_s$ neurons in each normalization head, and N neurons in the projection head. We made several tests with different configurations, and no improvement was observed by increasing this number. We use the rectified linear unit (ReLU) function as the activation function in the layers within the network backbone, whereas a linear function was employed in the multiple heads. We include a final concatenation layer to construct the action vector from the multiple heads.

The dimension of the input layer in the critic network equals $D_{\text{state}}^{\text{DCB}} + D_{\text{action}}^{\text{DCB}}$ since it takes both state and action vectors as inputs. As shown in Fig. 5, these vectors are first treated independently and later concatenated inside the network. The hidden layers are also fully connected with $2(D_{\text{state}}^{\text{DCB}} + D_{\text{action}}^{\text{DCB}})$ neurons. Similarly, we use the ReLU activation function in the hidden layers. The output layer dimension is one since this network aims to predict the value of the reward function for a given duple of state and action. Hence, we use the linear activation function at this output layer. Finally, we use the Adam optimizer in both networks since this algorithm has proven to be computationally efficient and robust for supervised and deep reinforcement learning problems [39], [52], [53].

V. RL-BASED JOINT OPTIMIZATION

In this section, we propose an RL-based approach to solving the optimization problem in (4). In the same way as for CB, we start addressing the selection of the state, action, and reward spaces. Next, we will review some important components in the DRL-based DDPG framework, and finally, we will derive the algorithmic solution.

A. DRL-DDPG: State, Action and Reward

As previously explained, the state, action, and reward spaces in RL must meet the MDP property, i.e., actions will affect the following states as well as the rewards. Besides, states must include information about all the aspects of the past agent-environment interactions that make a difference in the future.

Considering this idea, we have selected these elements in the following manner:

- The **state** vector \mathbf{s}_t is composed of the current values of all the precoders ($\mathbf{P}_k, \forall k$), the entries in the main diagonal of the IRS phase-shift matrix (θ), and entries in the channel response matrices (\mathbf{H}_{IB} and $\mathbf{H}_{\text{UI}k}, \forall k$). Thus, the state vector is constructed such that

$$\mathbf{s}_t = [\text{flatten}(\mathbf{P}_1), \dots, \text{flatten}(\mathbf{P}_K), \text{flatten}(\theta), \text{flatten}(\mathbf{H}_{\text{IB}}), \text{flatten}(\mathbf{H}_{\text{UI}1}), \dots, \text{flatten}(\mathbf{H}_{\text{UI}K})]. \quad (13)$$

Hence, the dimension of the state space vectors is $D_{\text{state}}^{\text{DRL}} = KN_t N_s + N + N_t N + KN_t N$. Recall that the state space is continuous since the entries of \mathbf{s}_t can take any complex value.

- The **action** vector \mathbf{a}_t comprises the matrices $\Delta\theta$ and $\Delta\mathbf{P}_k, \forall k$, whose entries stand for the variations in the values of θ and $\mathbf{P}_k, \forall k$, respectively. Hence, the dimension of the action space is $D_{\text{action}}^{\text{DRL}} = KN_t N_s + N$ and the action is constructed such that

$$\mathbf{a}_t = [\text{flatten}(\Delta\mathbf{P}_1), \dots, \text{flatten}(\Delta\mathbf{P}_K), \text{flatten}(\Delta\theta)]. \quad (14)$$

The action space is also continuous. However, its values are constrained to have imaginary and real parts within the real-valued interval $(-1, 1)$ to improve the system's stability.

- The **reward** r_t is determined as a function of the sum-rate, which is the metric we aim to maximize. In this DRL-DDPG formulation, rather than using the sum-rate value itself, we calculate r_t as the difference between the values after and before taking the action \mathbf{a}_t , i.e.,

$$r_t = \sum_{k=1}^K R_k(\mathbf{s}_t, \mathbf{a}_t) - \sum_{k=1}^K R_k(\mathbf{s}_{t-1}, \mathbf{a}_{t-1}). \quad (15)$$

In the way we have defined them, actions allow us to identify a correspondence between the current state and the action taken, and the corresponding next state—as expected in DRL-based formulations. Furthermore, the immediate reward depends on the state and the action, not only on the action.

The ANNs used for function approximation are sensitive to the scale of the features. If we use the sum-rate values as rewards, the difference between the scales of actions and rewards can be unfavorable to the learning and the stability of the consecutive time steps. By considering the reward to be the difference between the sum-rate values, the scales of state, action, and reward values remain similar.

B. DRL-DDPG: Framework Elements

Fig. 6 shows a representation of the DRL-DDPG agent. As observed, several of these elements were already used in the previously described DCB-DDPG framework. The only additional elements required for this framework are the **target networks**. In this case, the initial parameters of the actor and critic networks are copied into the target actor network

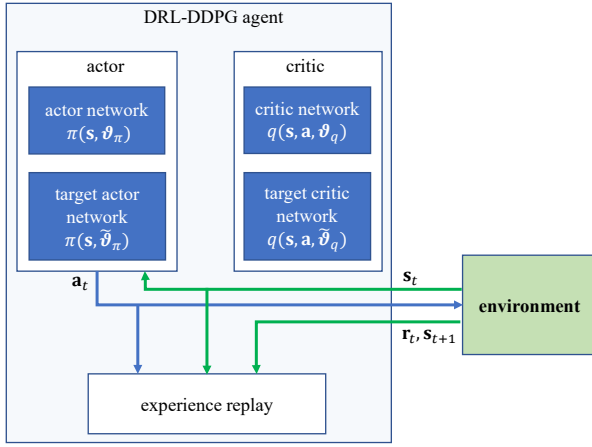


Fig. 6. Elements in the DRL-DDPG agent.

$\pi(s, \tilde{\vartheta}_\pi)$ and target critic network $q(s, \mathbf{a}, \tilde{\vartheta}_q)$, respectively. These target networks are used for calculating the target values, and their network parameters ($\tilde{\vartheta}_\pi$ and $\tilde{\vartheta}_q$) are updated in a soft-manner to track the original networks. Using these target networks makes it possible to set stable targets to improve the convergence of the critic network training. These target networks are not considered in the DCB approach since they are related to the estimations of future actions and rewards, which are relevant only to the DRL approach. Notice that in this case, we refer to the critic network as $q(s, \mathbf{a}, \vartheta_q)$ since it approximates the long-term reward function $q_\pi(\cdot)$.

C. DRL-DDPG Algorithm

In this section, we present the algorithm proposed to train the DRL-DDPG agent for computing the IRS phase-shift matrix (Θ) and the precoder (\mathbf{P}) that maximize the achievable sum-rate in (4). Algorithm 2 summarizes the steps to follow. Along the analysis, we focus on the main differences regarding the DCB-DDPG algorithm.

The actor and critic networks are created through the randomly initialized network parameters (ϑ_π and ϑ_q , respectively). In this DRL-DDPG approach, the target actor and target critic networks are also created as their exact copies. Hence, before the learning process starts, $\vartheta_\pi = \tilde{\vartheta}_\pi$ and $\vartheta_q = \tilde{\vartheta}_q$.

The interactions between the DRL-DDPG agent and the environment are split into E episodes with a finite number of time steps T since no terminal states can be defined in our problem. Dividing the training time steps into episodes enables the learning process to start from different initial states, thus allowing exploration over the entire state space. We consider setting $T = 20$ and varying the number of episodes according to the intended number of interactions.

The initial state s_0 is created as in (13) from the randomly initialized \mathbf{P} and Θ matrices, and the randomly generated channel response matrices ($\mathbf{H}_{\text{UI}k}, \forall k$ and \mathbf{H}_{IB}). Starting from that state, the agent takes actions according to the output of the actor network and an exploration noise $\mathbf{n}_e \sim \mathcal{N}_{\mathbb{C}}(0, \sigma_{\mathbf{n}_e}^2 \mathbf{I}_{D_{\text{action}}^{\text{DRL}}})$. As in the DCB-DDPG approach, the exploration noise guarantees a better exploration of the entire action space. Besides,

Algorithm 2 DRL-DDPG algorithm

```

1: Initialize:
2: set  $\pi(s, \vartheta_\pi)$  given  $\vartheta_\pi$ , and  $q(s, \mathbf{a}, \vartheta_q)$  given  $\vartheta_q$ 
3:  $\pi(s, \tilde{\vartheta}_\pi) \leftarrow \pi(s, \vartheta_\pi)$ , and  $q(s, \mathbf{a}, \tilde{\vartheta}_q) \leftarrow q(s, \mathbf{a}, \vartheta_q)$ 
4: create  $\mathcal{R}$ 
5: for  $e = 0, \dots, E - 1$  do:
6:   set  $s_0$  given random  $\mathbf{P} \in \mathcal{P}$  and  $\Theta \in \mathcal{D}$ 
7:   for  $t = 0, \dots, T - 1$  do:
8:     agent takes  $\mathbf{a}_t = \pi(s_t, \vartheta_\pi) + \mathbf{n}_e$ 
9:     environment returns  $r_t$  and  $s_{t+1}$ 
10:     $\mathcal{R}$  stores  $\mathcal{E}_t = (s_t, \mathbf{a}_t, r_t, s_{t+1})$ 
11:    if  $|\mathcal{R}| > |\mathcal{B}|$  :
12:      agent samples  $|\mathcal{B}|$  random experiences:
13:       $\mathcal{E}_i = (s_i, \mathbf{a}_i, r_i, s_{i+1}), i = 0, \dots, |\mathcal{B}| - 1$ 
14:      calculate  $\tilde{\mathbf{a}}_{i+1} = \pi(s_{i+1}, \tilde{\vartheta}_\pi)$ 
15:      calculate  $\tilde{y}_i = r_i + \gamma q(s_{i+1}, \tilde{\mathbf{a}}_{i+1}, \tilde{\vartheta}_q)$ 
16:      calculate the critic loss  $L_c$  by using (16)
17:      update  $\vartheta_q$  by the back-propagation of  $L_c$ 
18:      calculate the actor loss  $L_a$  by using (17)
19:      update  $\vartheta_\pi$  by the back-propagation of  $L_a$ 
20:       $\tilde{\vartheta}_\pi \leftarrow \tau \vartheta_\pi + (1 - \tau) \tilde{\vartheta}_\pi$ 
21:       $\tilde{\vartheta}_q \leftarrow \tau \vartheta_q + (1 - \tau) \tilde{\vartheta}_q$ 
22:    end if
23:  Obtain  $\mathbf{P}$  and  $\Theta$  by evaluating policy  $\pi(s, \vartheta_\pi)$ 

```

Output: trained actor network $\pi(s, \vartheta_\pi)$

in this DRL-DDPG approach, exploration noise also improves the exploration over the state space because actions affect the occurrence of the following states.

At each learning time step, the experience tuples in \mathcal{B} are employed to train the actor and critic networks as described in lines 12 to 19. As presented in [41], the target variables $\tilde{\mathbf{a}}_{i+1}$ and \tilde{y}_i are computed using the target actor network and the target critic network, respectively (lines 14 and 15). The factor γ in the computation of \tilde{y}_i stands for a discount rate and determines the relevance of future rewards. The critic loss L_c is next calculated as

$$L_c = \frac{1}{|\mathcal{B}|} \sum_i (\tilde{y}_i - q(s_i, \mathbf{a}_i, \vartheta_q))^2, \quad (16)$$

and the obtained values are back-propagated by using the critic network optimizer. This way, the critic network is trained to minimize L_c and approximate the actual behavior of the long-term reward function $q_\pi(\cdot)$.

On the other hand, the actor network is trained to predict the action that maximizes the output of the critic network. Hence, the actor loss L_a is calculated as

$$L_a = -\frac{1}{|\mathcal{B}|} \sum_i q(s_i, \pi(s_i, \vartheta_\pi), \vartheta_q), \quad (17)$$

and the obtained values are then utilized to update the parameters of the actor network.

The actor and critic target networks are updated as described in lines 20 and 21. Parameter τ defines the updating rate for the target networks. At the end of every learning time step, the obtained policy is evaluated by running an evaluation episode.

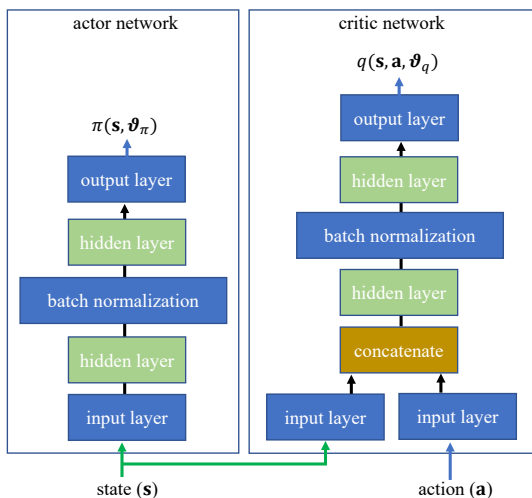


Fig. 7. DRL-DDPG actor and critic network structure.

We do not consider the exploration noise or update the networks during this stage. By following the actor network policy, the randomly initialized \mathbf{P} and Θ matrices are transformed to improve the sum-rate of the system.

D. DRL-DDPG: ANNs Structure

The structures of the ANNs we propose to use as actor and critic networks in DRL-DDPG are presented in Fig. 7. The input and output layer dimensions of the actor network equal $D_{\text{state}}^{\text{DRL}}$ and $D_{\text{action}}^{\text{DRL}}$, respectively. Two fully connected hidden layers with $2D_{\text{state}}$ neurons were considered for this network. As proposed in [39], [41], [54], we include batch normalization layers in both actor and critic networks to improve the learning performance. By inserting batch normalization layers, higher learning rates can be considered, and the network initialization values have a lower impact on the network performance. We used the ReLU function at both hidden layers, while the hyperbolic tangent (tanh) activation function was used at the output layer.

The critic network input and output dimensions equal $D_{\text{state}}^{\text{DRL}} + D_{\text{action}}^{\text{DRL}}$ and 1, respectively. The first layer, in this case, is a concatenation layer that transforms the state and action vectors into one. The inner structure of this network is similar to the actor but considers $2(D_{\text{state}}^{\text{DRL}} + D_{\text{action}}^{\text{DRL}})$ neurons at the hidden layers. The single-neuron output layer uses the linear activation function to predict the long-term reward value for the action and state inputs.

VI. CONVERGENCE AND COMPUTATIONAL COMPLEXITY

In this section, we analyze the convergence and computational complexity of the proposed approaches. We make a brief comparison to identify how learning and computing costs are affected by the framework structures and some configuration parameters.

A. Computational Complexity

In the following, we present a computational complexity analysis based on the number of required multiplications. As

stated in [55], this is a high-level metric for computational complexity evaluation because only multiplications are considered while other less time-consuming operations are disregarded. We consider two stages that occur at different moments and have different resource limitations. First, we address the computational complexity of the training stages (described in Algorithm 1 and Algorithm 2), which are expected to occur mostly offline. Second, we address the computational complexity of the operational stages where the already trained actor networks are used to predict the best IRS phase-shift and precoder matrices for the observed channel realizations.

Let us start by analyzing the computational complexity of the training stage in the DCB-DDPG approach (Algorithm 1). Since training can be performed for as many time steps as desired, we consider only the operations made within a time step in both approaches. The highest computational complexity in Algorithm 1 is in the operations with the critic network. When considering mini-batches of size $|\mathcal{B}|$, the computational complexity for ANNs with fully connected layers is bounded by $\mathcal{O}(|\mathcal{B}|\psi\zeta^2)$, where ψ is the number of hidden layers, and ζ is the number of neurons in the widest layer [55], [56]. We disregard the number of layers ψ because it is not related to the communication parameters, and its value is generally small compared to $|\mathcal{B}|$. The worst condition for ζ is considering the hidden layers of the critic network, which contain $2(D_{\text{state}}^{\text{DCB}} + D_{\text{action}}^{\text{DCB}})$ neurons. Hence, the computational complexity of this step, and the algorithm in general, is in the order of $\mathcal{O}(|\mathcal{B}|(K^2N_t^2N^2 + N_r^2N^2 + K^2N_t^2N_s^2 + N^2))$.

The computational complexity analysis for the training stage in DRL-DDPG is similar. In Algorithm 2, the highest complexity also corresponds to the calculations over the critic network. Hence, following the same reasoning, the computational complexity of this algorithm is bounded by $\mathcal{O}(|\mathcal{B}|(K^2N_t^2N^2 + N_r^2N^2 + 4K^2N_t^2N_s^2 + 4N^2))$. As observed, the computational cost of training in DRL-DDPG is higher due to the larger number of elements being considered in the states. Notice that, beyond the increment of the computational complexity, larger state vectors also imply larger storage capabilities needed for the ANN models and the replay buffer.

In the operational stage, the computational complexity is expected to be remarkably lower, which is desirable to fit the latency requirements of practical applications. In this stage, no exploration is performed, but experiences are stored for offline training. In the DCB-DDPG approach, the already trained actor network is used to predict the best IRS phase-shift and precoder matrices for a given channel realization in a single-step forward pass. As explained in Section IV, the widest layers in the DCB-DDPG actor network are in the backbone, with $2D_{\text{state}}^{\text{DCB}}$ neurons. Hence, the computational complexity of finding the best matrices in DCB-DDPG is in the order of $\mathcal{O}(K^2N_t^2N^2 + N_r^2N^2)$.

On the other hand, in the DRL-DDPG approach, the IRS phase-shift and precoder matrices are predicted through the sequence of T consecutive time steps that form an episode. In this case, the widest layers in the actor network for DRL-DDPG are composed of $D_{\text{state}}^{\text{DRL}}$ neurons. Hence, the computational complexity of the operational stage with DRL-DDPG is in the order of

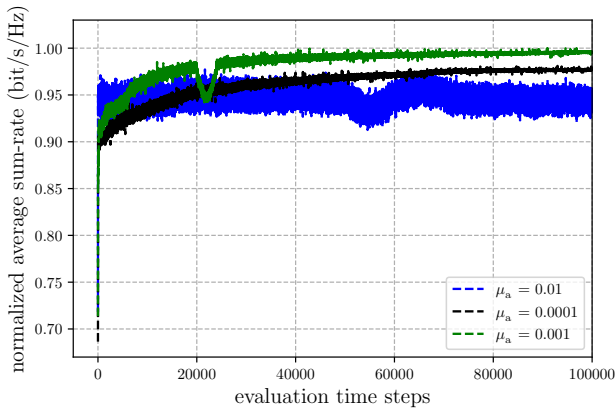


Fig. 8. Normalized average sum-rate (bit/s/Hz) vs evaluation time steps with the DCB-DDPG framework.

$$\mathcal{O}(T(K^2 N_t^2 N^2 + N_r^2 N^2 + K^2 N_t^2 N_s^2 + N^2)).$$

As shown before, the DCB-DDPG approach outperforms the DRL-DDPG scheme in terms of computational complexity. By presenting simpler implementations of the ANN-based function approximations, and a simpler algorithm, DCB-DDPG stands as a more suitable alternative to handle the joint optimization problem, especially in applications with stringent latency requirements or computing resource constraints.

B. Convergence

As explained in [34], convergence cannot be guaranteed in algorithms like those proposed in DCB-DDPG and DRL-DDPG. These algorithms combine three crucial features, which are known as *the deadly triad*: function approximation, bootstrapping (i.e., using previous estimates for training), and off-policy training. They are called this way because instability and divergence arise when these features are jointly employed. However, none of these features can be given up because of their relevance. For this reason, the convergence analysis must be empirically addressed by considering the specific conditions of our frameworks and system model.

As in [39], [40], [42], [56], [57], we first propose a convergence analysis based on the reward function. Fig. 8 and Fig. 9 show the normalized average sum-rate values computed at the evaluation time steps in DCB-DDPG and the evaluation episodes in DRL-DDPG, respectively. At each evaluation stage, the average sum-rate values are computed over 100 channel realizations different from those visited during training. Normalization is performed by considering the overall highest average sum-rate value as the normalization factor to allow a fair comparison. These figures show that convergence in terms of sum-rate can be empirically illustrated for both approaches in the considered setup. The obtained values steadily improve and converge to their best solutions when the learning rates of the actor network optimizers (μ_a) are equal to 0.001. Thus, the proper behavior of the actor networks is demonstrated, since they continuously improve on learning the best precoder and the IRS phase-shift matrix for each channel realization.

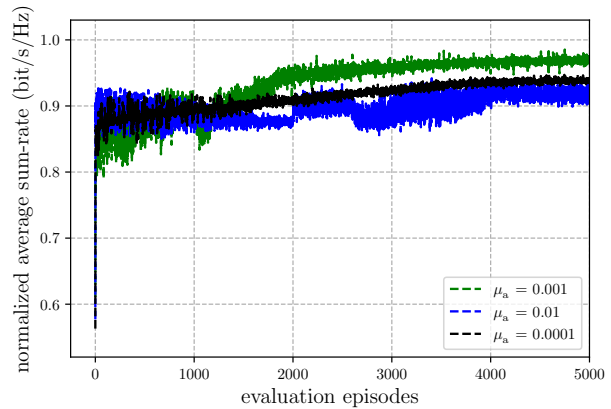


Fig. 9. Normalized average sum-rate (bit/s/Hz) vs evaluation episodes with the DRL-DDPG framework.




Fig. 10. Normalized critic loss vs training time steps with the DCB-DDPG framework.

Additionally, we have performed a second convergence analysis based on the critic loss values (L_c). The better the critic networks approximate the instantaneous reward function $r_t(\cdot)$ in DCB-DDPG and the long-term reward function $q_\pi(\cdot)$ in DRL-DDPG, the smaller the values of L_c are. Fig. 10 and Fig. 11 show the normalized critic loss values in the DCB-DDPG and the DRL-DDPG approaches, respectively. As in the previous analysis, the convergence of both algorithms can be empirically shown for at least one of the learning rates of the critic network optimizer ($\mu_c = 0.001$). Notice that, in both convergence analyses, higher learning rate values ($\mu_a = \mu_c = 0.01$) lead to lower stability. On the other hand, lower learning rates ($\mu_a = \mu_c = 0.0001$) improve stability but sacrifice convergence speed.

As observed in the previous figures, the proposed DCB-DDPG framework is more robust regarding sum-rate maximization and critic loss minimization. In both cases, DCB-DDPG converges to better solutions while keeping lower variance values. In this sense, convergence speed and stability are mainly related to the complexity of the learning and evaluation stages. In DCB-DDPG, single-time step predictions are used to find the best precoder and IRS phase-shift matrices. Hence,

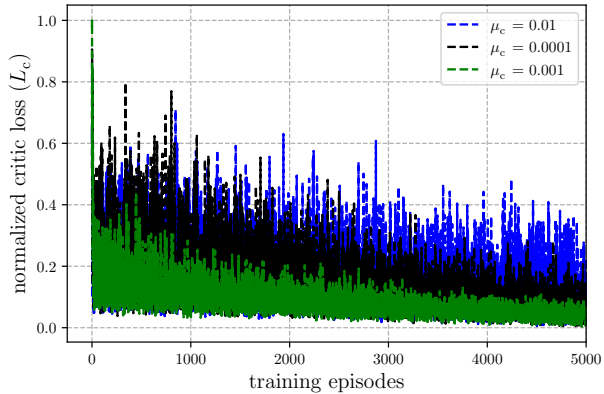


Fig. 11. Normalized critic loss vs training episodes with the DRL-DDPG framework.

either good or bad, predictions in DCB-DDPG only affect the instantaneous reward and not the future states or rewards. On the other hand, matrices in DRL-DDPG are predicted through the sequence of consecutive time steps that form an episode. Hence, prediction errors might propagate and increase from one time step to the next, thus reducing the system stability.

VII. SIMULATION RESULTS

In this section, we present results that validate the use of the DCB-DDPG and DRL-DDPG frameworks to optimize the IRS phase-shift and precoder matrices in the uplink of an IRS-assisted MS MU-MIMO system, as the one described in Section III. We consider the setups presented in Sections IV and V, and follow the steps described in Algorithm 1 and Algorithm 2, respectively.

Table III shows the configuration parameters considered for the computer experiments. For fairness, we have selected the number of episodes in DRL-DDPG and the time steps in both proposals so that the training times remain similar for both approaches. In both cases, the mini-batches with 16 entries constitute an adequate trade-off between learning performance and complexity. The learning rate values in both proposals and the update and discount factors in DRL-DDPG were obtained experimentally through a grid search approach to provide the best performance in terms of the system's achievable sum-rate. The exploration noise variance was set to 0.05 in both approaches to guarantee a fair trade-off between exploration and exploitation over the action space. This value is also in concordance with the one proposed in [41]. As explained, exploration over state space is also guaranteed by the randomness of channel realizations and the episodes in the DRL-DDPG approach.

We have selected two model-driven approaches as benchmarks. The first scheme, termed Alternating-ProG, uses an alternating minimization projected gradient algorithm to optimize the IRS and precoder matrices [58, Algorithm 5]. This solution has been shown to perform better than other strategies proposed in the literature for the joint design of the IRS and precoding matrices. We have also included a baseline strategy

TABLE III
CONFIGURATION PARAMETERS

DCB-DDPG parameters		
Parameter	Description	Value
T	number of training time steps	100000
$ \mathcal{B} $	mini-batch size	16
$ \mathcal{R} _{\max}$	max. experience replay buffer size	100000
μ_a	actor network optimizer learning rate	0.001
μ_c	critic network optimizer learning rate	0.001
$\sigma_{n_e}^2$	exploration noise variance	0.05
DRL-DDPG parameters		
Parameter	Description	Value
E	number of episodes	5000
T	number of time steps per episode	20
$ \mathcal{B} $	mini-batch size	16
$ \mathcal{R} _{\max}$	max. experience replay buffer size	100000
μ_a	actor network optimizer learning rate	0.001
μ_c	critic network optimizer learning rate	0.001
τ	target network updating rate	0.005
γ	discount factor	0.99
$\sigma_{n_e}^2$	exploration noise variance	0.05




Fig. 12. Sum-rate (bit/s/Hz) vs SNR (dB) for $K = 10$, $N_s = 2$, $N_t = 2$, $N_r = 30$, $N = 50$.

termed RandomIRS-MRT to provide a lower-bound reference. In the RandomIRS-MRT scheme, the IRS phase-shift matrix is randomly selected from the set of feasible matrices \mathcal{D} , and the precoders are computed according to the maximum ratio transmitter (MRT) criterion.

Fig. 12 shows the achievable sum-rates obtained with the proposed DCB-DDPG and DRL-DDPG approaches and the two benchmarks for a range of SNR values between -15 dB and 15 dB. The results in Fig. 12 were obtained by considering a setup with $K = 10$ users employing $N_t = 2$ antennas to send $N_s = 2$ streams each, an IRS with $N = 50$ scattering elements, and a BS with $N_r = 30$ receiving antennas.

In these simulation conditions, the Alternating-ProG benchmark outperforms our proposals, although the difference concerning the DCB-DDPG scheme is minimal. This is an expected result that an important remark can explain. As stated in [15], the performance of linear algorithms is near optimal when $N_r \gg KN_s$, as occurs in this configuration. Although the necessary alternating procedure imposes certain limitations on its performance, we can say that the communication configuration is favorable for the benchmark approach. The

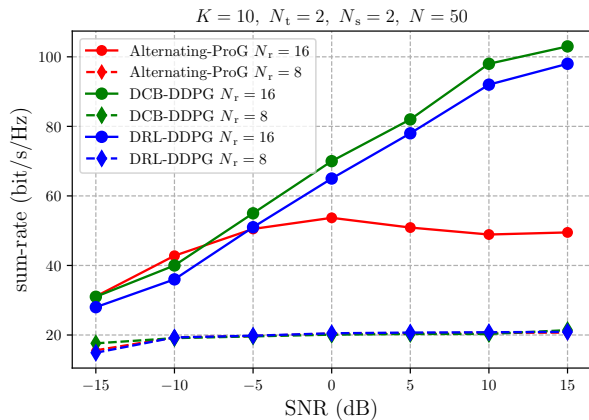


Fig. 13. Sum-rate (bit/s/Hz) vs SNR (dB) for $N_r \in \{8, 16\}$, $K = 10$, $N_t = 2$, $N_s = 2$, $N = 50$.

performance of DRL-DDPG in this scenario is also good, although it is outperformed by the DCB-DDPG algorithm. Since we constrained the training times of both approaches to be equal, the higher complexity of DRL-DDPG and better suitability of the DCB formulation for this problem led to the observed gap in their performances. On the other hand, both proposals significantly outperform the RandomIRS-MRT baseline strategy.

We next considered the same setup as in the previous experiment and conducted simulations for two different numbers of BS receiving antennas, namely $N_r \in \{8, 16\}$. The results are shown in Fig. 13. Note that for $N_r = 16$, the total number of transmitted data streams $KN_s = 20$ is slightly higher than the number of receiving antennas N_r . In the low SNR regime (SNR < -5 dB), the behavior of the sum-rate values is similar to that observed in Fig. 12. The DCB-DDPG and Alternating-ProG schemes behave similarly and slightly outperform DRL-DDPG. This result is reasonable because, in the low SNR regime, the system performance is more sensitive to channel noise than to multi-user interference. However, when SNR increases, Alternating-ProG is unable to handle the multi-user interference, and the achievable sum-rate decreases or remains constant. On the other hand, both DCB-DDPG and DRL-DDPG manage the interference more appropriately, and the sum-rate increases almost linearly with the SNR values. For example, the achievable sum-rate obtained with DCB-DDPG when SNR = 15 dB is more than 40 bits/s/Hz higher than that provided by Alternating-ProG. In addition, DCB-DDPG outperforms DRL-DDPG by around 5 bits/s/Hz in all the SNR regimes.

Fig. 13 also shows the results obtained when $N_r = 8$. This is a highly demanding setup since the total of transmitted data streams $KN_s = 20$ is higher than twice the number of receiving antennas. Under this condition, none of the algorithms is able to tackle the multi-user interference, and a scheduling stage should be considered.

We conducted another set of experiments to analyze the performance impact of the relationship between the number of receiving antennas N_r and the number of transmitted streams




Fig. 14. Sum-rate (bit/s/Hz) vs N_r for SNR = 10 dB, $K = 10$, $N_t = 2$, $N_s = 2$, $N = 50$.

KN_s . We again assumed a setup with $K = 10$, $N_s = 2$, $N_t = 2$ and $N = 50$. We fixed SNR = 10 dB and considered N_r values that range from 8 to 30. Fig. 14 shows the obtained results. In the following, we analyzed them by considering three working regimes: $N_r/KN_s \ll 1$, $N_r/KN_s < 1$, and $N_r/KN_s \geq 1$.

The first regime with $N_r/KN_s \ll 1$ represents a critical scenario where, as discussed previously, none of the analyzed algorithms perform well, and their sum-rates degrade remarkably. In the third regime with $N_r/KN_s \geq 1$, all the algorithms perform reasonably well, although DCB-DDPG and Alternating-ProG outperform DRL-DDPG. It is also worth noting that the DCB-DDPG scheme nearly achieves the same performance as the model-driven benchmark in this regime. This behavior matches that observed in Fig. 12. The second regime with $N_r/KN_s < 1$ is an intermediate scenario where DCB-DDPG and DRL-DDPG clearly outperform Alternating-ProG. When moving into this regime, Alternating-ProG cannot properly handle the interference among the users. The sum-rate values slowly increase when considering more receiving antennas, and no major performance improvement is observed until the third regime condition is met. On the other hand, DCB-DDPG and DRL-DDPG manage the interference more efficiently, and the sum-rate values increase almost linearly when considering more receiving antennas. When N_r decreases from 20 to 18, the Alternating-ProG sum-rate value degrades near 75 bit/s/Hz, while this value only degrades around 25 bit/s/Hz for DCB-DDPG and DRL-DDPG. The results show that the sum-rate values achieved with the proposed schemes are at least 25 bit/s/Hz higher when considering N_r values between 12 and 18.

As explained in Section II, formulating the joint optimization problem as RL or CB with discrete actions is not feasible. The number of possible actions becomes intractable even for low-complexity setups. The number of possible actions—and the size of the output layer of the ANN used for the policy function approximation—would be as large as $\approx 10^{15}$, only by addressing the optimization of an IRS with 50 scattering elements and two possible phase values. This

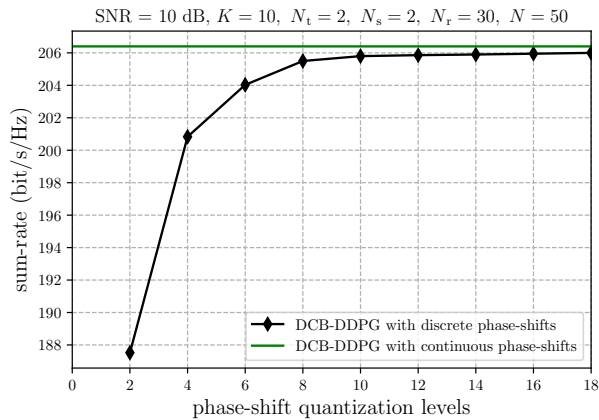


Fig. 15. Sum-rate (bit/s/Hz) vs IRS phase-shift matrix quantization levels.

number is several times larger if we also consider discrete-valued precoding vectors. However, to meet the requirements of practical implementations, we have analyzed the use of an additional discretization stage to be applied to the already computed optimal continuous-valued IRS matrices, instead of considering discrete formulations of the problem.

Fig. 15 shows the sum-rate when the discretization stage is applied to the IRS phase-shift matrix obtained for a configuration with $K = 10$, $N_t = 2$, $N_s = 2$, $N_r = 30$, $N = 50$ and $\text{SNR} = 10$ dB. We use the continuous-valued IRS phase-shift matrix obtained by the DCB-DDPG framework as a baseline. The figure shows that the sum-rate degradation caused by discretizing the IRS phase-shifts is lower than 20 bit/s/Hz, even considering only two possible phases for the IRS scattering elements. Besides, this degradation effect steeply decreases when more phase-shift quantization levels are considered. As observed, we reach almost the same performance with only eight quantization levels. Comparable results have been obtained for other system configurations. In addition, a similar codebook-based approach could be considered for the precoders. Hence, continuous-valued formulations are a good choice for the system design, even if discrete-valued matrices are the final objective.

VIII. CONCLUSIONS

We have investigated two different approaches for the joint optimization of the IRS phase-shift matrix and the precoders in the uplink of an IRS-assisted MS MU-MIMO system. The first approach is a CB formulation with continuous state and action spaces. To handle this continuous formulation, we have developed an actor-critic framework called DCB-DDPG. The second approach is a DRL formulation of the optimization problem. By considering the pros and cons of related works, we developed an approach to the DDPG framework with an alternative configuration of the state and action spaces, such that entries in the current matrices are the states and the variations in these values are the actions. A deep analysis of both proposed frameworks has been made, and our main findings can be summarized as follows:

- CB and RL enable an adequate formulation of the joint optimization, while DCB-DDPG and DRL-DDPG frameworks are effective methods to solve the considered optimization problem.
- The convergence of both proposed frameworks in terms of reward maximization and critic loss minimization is empirically illustrated. However, DCB-DDPG shows better stability and convergence speed, which can be related to its lower computational complexity and better suitability of the problem formulation.
- The performance of both proposed frameworks regarding sum-rate maximization is similar to state-of-the-art heuristic algorithms when $N_r \geq KN_s$. In contrast, the proposed schemes handle the multi-user interference more efficiently when $N_r < KN_s$.
- Simulation results show that our continuous-valued formulations are valuable even when discrete-valued matrices are required as final outcomes.

ACKNOWLEDGMENTS

This work has been supported by grants ED431C 2020/15 and ED431G 2019/01 (to support the Centro de Investigación de Galicia “CITIC”) funded by Xunta de Galicia and ERDF Galicia 2014-2020; and by grants PID2019-104958RB-C42 (ADELE) and BES-2017-081955 funded by MCIN/AEI/10.13039/501100011033.

REFERENCES

- [1] Y.-C. Liang, R. Long, Q. Zhang, J. Chen, H. V. Cheng, and H. Guo, “Large Intelligent Surface/Antennas (LISA): Making Reflective Radios Smart,” *arXiv:1906.06578 [cs, math]*, Jun. 2019. [Online]. Available: <http://arxiv.org/abs/1906.06578>
- [2] S. Gong, X. Lu, D. T. Hoang, D. Niyato, L. Shu, D. I. Kim, and Y.-C. Liang, “Toward Smart Wireless Communications via Intelligent Reflecting Surfaces: A Contemporary Survey,” *IEEE Commun. Surveys Tuts.*, vol. 22, no. 4, pp. 2283–2314, 2020.
- [3] X. Guan, Q. Wu, and R. Zhang, “Joint Power Control and Passive Beamforming in IRS-Assisted Spectrum Sharing,” *IEEE Commun. Lett.*, vol. 24, no. 7, pp. 1553–1557, Jul. 2020.
- [4] E. Castañeda, A. Silva, A. Gameiro, and M. Kountouris, “An Overview on Resource Allocation Techniques for Multi-User MIMO Systems,” *IEEE Commun. Surveys Tuts.*, vol. 19, no. 1, pp. 239–284, 2017.
- [5] J. Navarro-Ortiz, P. Romero-Diaz, S. Sendra, P. Ameigeiras, J. J. Ramos-Munoz, and J. M. Lopez-Soler, “A Survey on 5G Usage Scenarios and Traffic Models,” *IEEE Commun. Surveys Tuts.*, vol. 22, no. 2, pp. 905–929, 2020.
- [6] L. Zhao, H. Zhao, K. Zheng, and W. Xiang, *Massive MIMO in 5G Networks: Selected Applications*, ser. SpringerBriefs in Electrical and Computer Engineering. Cham: Springer International Publishing, 2018. [Online]. Available: <http://link.springer.com/10.1007/978-3-319-68409-3>
- [7] D. Pérez-Adán, O. Fresnedo, J. P. González-Coma, and L. Castedo, “Intelligent Reflective Surfaces for Wireless Networks: An Overview of Applications, Approached Issues, and Open Problems,” *Electronics*, vol. 10, no. 19, p. 2345, 2021. [Online]. Available: <https://www.mdpi.com/2079-9292/10/19/2345>
- [8] C. Huang, A. Zappone, G. C. Alexandropoulos, M. Debbah, and C. Yuen, “Reconfigurable Intelligent Surfaces for Energy Efficiency in Wireless Communication,” *IEEE Trans. Wireless Commun.*, vol. 18, no. 8, pp. 4157–4170, Aug. 2019.
- [9] H. Guo, Y.-C. Liang, J. Chen, and E. G. Larsson, “Weighted Sum-Rate Maximization for Reconfigurable Intelligent Surface Aided Wireless Networks,” *IEEE Trans. Wireless Commun.*, vol. 19, no. 5, pp. 3064–3076, May 2020.
- [10] A. Al-Hilo, M. Samir, M. Elhattab, C. Assi, and S. Sharafeddine, “Reconfigurable Intelligent Surface Enabled Vehicular Communication: Joint User Scheduling and Passive Beamforming,” *IEEE Trans. Veh. Technol.*, vol. 71, no. 3, pp. 2333–2345, Mar. 2022.

- [11] G. Geraci, A. Garcia-Rodriguez, M. M. Azari, A. Lozano, M. Mezzavilla, S. Chatzinotas, Y. Chen, S. Rangan, and M. Di Renzo, "What Will the Future of UAV Cellular Communications Be? A Flight from 5G to 6G," *arXiv:2105.04842 [cs, eess, math]*, May 2021. [Online]. Available: <http://arxiv.org/abs/2105.04842>
- [12] Z. Wei, Y. Cai, Z. Sun, D. W. K. Ng, J. Yuan, M. Zhou, and L. Sun, "Sum-Rate Maximization for IRS-Assisted UAV OFDMA Communication Systems," *IEEE Trans. Wireless Commun.*, vol. 20, no. 4, pp. 2530–2550, Apr. 2021.
- [13] Z. Chen, X. Ma, C. Han, and Q. Wen, "Towards intelligent reflecting surface empowered 6G terahertz communications: A survey," *China Commun.*, vol. 18, no. 5, pp. 93–119, May 2021.
- [14] X. Ma, Z. Chen, W. Chen, Z. Li, Y. Chi, C. Han, and S. Li, "Joint Channel Estimation and Data Rate Maximization for Intelligent Reflecting Surface Assisted Terahertz MIMO Communication Systems," *IEEE Access*, vol. 8, pp. 99 565–99 581, 2020.
- [15] E. Björnson, J. Hoydis, and L. Sanguinetti, "Massive MIMO Networks: Spectral, Energy, and Hardware Efficiency," *FNT in Signal Processing*, vol. 11, no. 3–4, pp. 154–655, 2017. [Online]. Available: <http://www.nowpublishers.com/article/Details/SIG-093>
- [16] M. Hong and Z.-Q. Luo, "Signal Processing and Optimal Resource Allocation for the Interference Channel," *arXiv:1206.5144 [cs, math]*, Jun. 2012. [Online]. Available: <http://arxiv.org/abs/1206.5144>
- [17] W. Wang, X. Liu, J. Tang, N. Zhao, Y. Chen, Z. Ding, and X. Wang, "Beamforming and Jamming Optimization for IRS-Aided Secure NOMA Networks," *IEEE Trans. Wireless Commun.*, pp. 1–1, 2021.
- [18] H. Ur Rehman, F. Bellili, A. Mezghani, and E. Hossain, "Joint Active and Passive Beamforming Design for IRS-Assisted Multi-User MIMO Systems: A VAMP-Based Approach," *IEEE Trans. Commun.*, vol. 69, no. 10, pp. 6734–6749, Oct. 2021.
- [19] D. Chang, H. Jiang, J. Zhou, H. Zhang, and M. Mukherjee, "Capacity Optimization using Augmented Lagrange Method in Intelligent Reflecting Surface-based MIMO Communication systems," *China Commun.*, vol. 17, no. 12, pp. 123–138, Dec. 2020.
- [20] K. Zhang, C. Liu, H. Wang, and Y. Song, "An IRS-Aided mmWave Massive MIMO Systems Based on Genetic Algorithm," in *Proc. 20th IEEE Int. Conf. Commun. Technol. (ICCT)*, Oct. 2020, pp. 288–293, iSSN: 2576-7828.
- [21] V. P. Tuan and I. P. Hong, "Secrecy Performance Analysis and Optimization of Intelligent Reflecting Surface-Aided Indoor Wireless Communications," *IEEE Access*, vol. 8, pp. 109 440–109 452, 2020.
- [22] S. Huang, Y. Ye, M. Xiao, H. V. Poor, and M. Skoglund, "Decentralized Beamforming Design for Intelligent Reflecting Surface-Enhanced Cell-Free Networks," *IEEE Wireless Commun. Lett.*, vol. 10, no. 3, pp. 673–677, Mar. 2021.
- [23] S. Liu, Z. Gao, J. Zhang, M. D. Renzo, and M.-S. Alouini, "Deep Denoising Neural Network Assisted Compressive Channel Estimation for mmWave Intelligent Reflecting Surfaces," *IEEE Trans. Veh. Technol.*, vol. 69, no. 8, pp. 9223–9228, Aug. 2020.
- [24] A. M. Elbir, A. Papazafeiropoulos, P. Kourtessis, and S. Chatzinotas, "Deep Channel Learning for Large Intelligent Surfaces Aided mm-Wave Massive MIMO Systems," *IEEE Wireless Commun. Lett.*, vol. 9, no. 9, pp. 1447–1451, Sep. 2020.
- [25] S. Yin, Y. Li, Y. Tian, and R. Yu, "Intelligent Reflecting Surface Enhanced Wireless Communications with Deep-learning-based Channel Prediction," *IEEE Trans. Veh. Technol.*, pp. 1–1, 2021.
- [26] W. Xia, G. Zheng, Y. Zhu, J. Zhang, J. Wang, and A. P. Petropulu, "A Deep Learning Framework for Optimization of MISO Downlink Beamforming," *IEEE Trans. Commun.*, vol. 68, no. 3, pp. 1866–1880, Mar. 2020.
- [27] Y. Ge and J. Fan, "Beamforming Optimization for Intelligent Reflecting Surface Assisted MISO: A Deep Transfer Learning Approach," *IEEE Trans. Veh. Technol.*, vol. 70, no. 4, pp. 3902–3907, Apr. 2021.
- [28] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. The MIT Press, 2016. [Online]. Available: <http://www.deeplearningbook.org>
- [29] N. Zhao, Y.-C. Liang, D. Niyato, Y. Pei, M. Wu, and Y. Jiang, "Deep Reinforcement Learning for User Association and Resource Allocation in Heterogeneous Cellular Networks," *IEEE Trans. Wireless Commun.*, vol. 18, no. 11, pp. 5141–5152, Nov. 2019.
- [30] C. Huang, G. Chen, Y. Gong, M. Wen, and J. A. Chambers, "Deep Reinforcement Learning-Based Relay Selection in Intelligent Reflecting Surface Assisted Cooperative Networks," *IEEE Wireless Commun. Lett.*, vol. 10, no. 5, pp. 1036–1040, May 2021.
- [31] H. Yang, Z. Xiong, J. Zhao, D. Niyato, Q. Wu, H. V. Poor, and M. Tornatore, "Intelligent Reflecting Surface Assisted Anti-Jamming Communications: A Fast Reinforcement Learning Approach," *IEEE Trans. Wireless Commun.*, vol. 20, no. 3, pp. 1963–1974, Mar. 2021.
- [32] N. Nurani Krishnan, E. Torkildson, N. B. Mandayam, D. Raychaudhuri, E.-H. Rantala, and K. Doppler, "Optimizing Throughput Performance in Distributed MIMO Wi-Fi Networks Using Deep Reinforcement Learning," *IEEE Trans. on Cogn. Commun. Netw.*, vol. 6, no. 1, pp. 135–150, Mar. 2020.
- [33] L. Zhang and Y.-C. Liang, "Deep Reinforcement Learning for Multi-Agent Power Control in Heterogeneous Networks," *IEEE Trans. Wireless Commun.*, vol. 20, no. 4, pp. 2551–2564, Apr. 2021.
- [34] R. S. Sutton and A. G. Barto, *Reinforcement learning: an introduction*, second edition ed., ser. Adaptive computation and machine learning series. Cambridge, Massachusetts: The MIT Press, 2018.
- [35] I. Ismath, S. Ali, N. Rajatheva, and M. Latva-Aho, "Deep Contextual Bandits for Fast Neighbor-Aided Initial Access in mmWave Cell-Free Networks," *IEEE Wireless Commun. Lett.*, vol. 10, no. 12, pp. 2752–2756, Dec. 2021.
- [36] J. Zhang and C. Masouros, "Beam Drift in Millimeter Wave Links: Beamwidth Tradeoffs and Learning Based Optimization," *IEEE Trans. Commun.*, vol. 69, no. 10, pp. 6661–6674, Oct. 2021.
- [37] E. M. Mohamed, S. Hashima, K. Hatano, M. M. Fouda, and Z. M. Fadlullah, "Sleeping Contextual/Non-Contextual Thompson Sampling MAB for mmWave D2D Two-Hop Relay Probing," *IEEE Trans. Veh. Technol.*, vol. 70, no. 11, pp. 12 101–12 112, Nov. 2021.
- [38] F. Pan, Q. Cai, P. Tang, F. Zhuang, and Q. He, "Policy gradients for contextual recommendations," in *The World Wide Web Conference on - WWW '19*. ACM Press, 2019. [Online]. Available: <https://doi.org/10.1145%2F3308558.3313616>
- [39] C. Huang, R. Mo, and C. Yuen, "Reconfigurable Intelligent Surface Assisted Multiuser MISO Systems Exploiting Deep Reinforcement Learning," *IEEE J. Sel. Areas Commun.*, vol. 38, no. 8, pp. 1839–1850, Aug. 2020.
- [40] K. Stylianopoulos, G. Alexandropoulos, C. Huang, C. Yuen, M. Bennis, and a. M. Debbah, "Deep Contextual Bandits for Orchestrating Multi-User MISO Systems with Multiple RISs," *arXiv:2202.08194 [cs, eess, math]*, Feb. 2022. [Online]. Available: <http://arxiv.org/abs/2202.08194>
- [41] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous Control with Deep Reinforcement Learning," *arXiv:1509.02971 [cs, stat]*, Jul. 2019. [Online]. Available: <http://arxiv.org/abs/1509.02971>
- [42] J. Xu and B. Ai, "Experience-Driven Power Allocation Using Multi-Agent Deep Reinforcement Learning for Millimeter-Wave High-Speed Railway Systems," *IEEE Trans. Intell. Transp. Syst.*, pp. 1–11, 2021.
- [43] D. Pereira-Ruisanchez, O. Fresnedo, D. Perez-Adan, and L. Castedo, "Joint Optimization of IRS-assisted MU-MIMO Communication Systems through a DRL-based Twin Delayed DDPG Approach," in *2022 IEEE International Symposium on Broadband Multimedia Systems and Broadcasting (BMSB)*, Jun. 2022, pp. 1–6, iSSN: 2155-5052.
- [44] H. Albinsaid, K. Singh, S. Biswas, and C.-P. Li, "Multi-agent Reinforcement Learning Based Distributed Dynamic Spectrum Access," *IEEE Trans. on Cogn. Commun. Netw.*, pp. 1–1, 2021.
- [45] X. Guo, Y. Chen, and Y. Wang, "Learning-based Robust and Secure Transmission for Reconfigurable Intelligent Surface Aided Millimeter Wave UAV Communications," *IEEE Wireless Commun. Lett.*, pp. 1–1, 2021, conference Name: IEEE Wireless Communications Letters.
- [46] D. Silver, G. Lever, N. Heess, T. Degris, D. Wierstra, and M. Riedmiller, "Deterministic Policy Gradient Algorithms," in *Proc. 31st Int. Conf. Mach. Learn. (ICML)*. PMLR, Jan. 2014, pp. 387–395. [Online]. Available: <https://proceedings.mlr.press/v32/silver14.html>
- [47] A. Zai and B. Brown, *Deep Reinforcement Learning in Action*. New York, USA: Manning Publications Co., 2020.
- [48] Q. Wu and R. Zhang, "Beamforming Optimization for Wireless Network Aided by Intelligent Reflecting Surface With Discrete Phase Shifts," *IEEE Trans. Commun.*, vol. 68, no. 3, pp. 1838–1851, Mar. 2020.
- [49] Z. Wang, L. Liu, and S. Cui, "Channel Estimation for Intelligent Reflecting Surface Assisted Multiuser Communications: Framework, Algorithms, and Analysis," *IEEE Trans. Wireless Commun.*, vol. 19, no. 10, pp. 6607–6620, Oct. 2020.
- [50] M. Joham, H. Gao, and W. Utschick, "Estimation Of Channels In Systems With Intelligent Reflecting Surfaces," in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2022, pp. 5368–5372.
- [51] C. You, B. Zheng, and R. Zhang, "Channel Estimation and Passive Beamforming for Intelligent Reflecting Surface: Discrete Phase Shift and Progressive Refinement," *IEEE J. Sel. Areas Commun.*, vol. 38, no. 11, pp. 2604–2620, Nov. 2020.

- [52] K. Feng, Q. Wang, X. Li, and C.-K. Wen, "Deep Reinforcement Learning Based Intelligent Reflecting Surface Optimization for MISO Communication Systems," *IEEE Wireless Commun. Lett.*, vol. 9, no. 5, pp. 745–749, May 2020.
- [53] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," *arXiv:1412.6980 [cs]*, Jan. 2017. [Online]. Available: <http://arxiv.org/abs/1412.6980>
- [54] S. Ioffe and C. Szegedy, "Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift," *arXiv:1502.03167 [cs]*, Mar. 2015. [Online]. Available: <http://arxiv.org/abs/1502.03167>
- [55] P. J. Freire, S. Srivallapanondh, A. Napoli, J. E. Prilepsky, and S. K. Turitsyn, "Computational Complexity Evaluation of Neural Network Applications in Signal Processing," Jun. 2022. [Online]. Available: <http://arxiv.org/abs/2206.12191>
- [56] Y. Zhao, I. G. Niemegeers, and S. M. H. De Groot, "Dynamic Power Allocation for Cell-Free Massive MIMO: Deep Reinforcement Learning Methods," *IEEE Access*, vol. 9, pp. 102 953–102 965, 2021.
- [57] H. Yang, Z. Xiong, J. Zhao, D. Niyato, L. Xiao, and Q. Wu, "Deep Reinforcement Learning-Based Intelligent Reflecting Surface for Secure Wireless Communications," *IEEE Trans. Wireless Commun.*, vol. 20, no. 1, pp. 375–388, Jan. 2021.
- [58] N. S. Perović, L.-N. Tran, M. Di Renzo, and M. F. Flanagan, "On the Maximum Achievable Sum-rate of the RIS-aided MIMO Broadcast Channel," *arXiv:2110.01700 [cs, math]*, Oct. 2021, arXiv: 2110.01700. [Online]. Available: <http://arxiv.org/abs/2110.01700>



Daríel Pereira-Ruisánchez (Student Member, IEEE) received the B.S. degree in Telecommunications and Electronics Engineering from the Technological University of Havana José Antonio Echeverría (CUJAE), Cuba, in 2017. He is currently pursuing the Ph.D. degree at the University of A Coruña (UDC), Spain, where he has been with the CITIC, Centre for Information and Communications Technology Research and the Group of Electronic Technology and Communications (GTEC), since 2021. He was awarded a Marie Skłodowska-Curie predoctoral fellowship within the 3-i ICT Ph.D. Programme hosted by CITIC. His research interests include wireless communication systems, deep learning-based signal processing techniques, and broadcast applications.



Óscar Fresnedo (Member, IEEE) received the Computer Engineering and the PhD in Computer Engineering from the University of A Coruña, Spain, in 2007 and 2014, respectively. Since 2007 he is with the Group of Electronic Technology and Communications (GTEC) in the Department of Electronics and Systems at the University of A Coruña, where he had the benefit of an FPI scholarship granted by the Spanish Government from 2008 to 2012. His main research interests are in the design of coding schemes, analog joint source-channel coding, multi-user communications, and image processing. He has published 20 papers in international technical journals as well as more than 30 papers in relevant international conferences and workshops in the area of communications and signal processing. He has participated as a research member in more than 20 research projects and contracts granted by regional, national and European administrations. Dr. Oscar Fresnedo has also received the Best Student Paper Award at the 14th IEEE International Workshop on Signal Processing Advances in Wireless Communications (SPAWC), Darmstadt, 2013.



Darian Pérez-Adán received the B.S. degree in telecommunications and electronics engineering from the Technological University of Havana José Antonio Echeverría (CUJAE), Cuba, in 2017, and the Ph.D. degree in computer engineering from the University of A Coruña (UDC), Spain, in 2022. Since 2018, he has been with the Group of Electronic Technology and Communications, UDC. He was awarded a predoctoral scholarship granted by the Spanish Government. He was a visiting researcher with the Associate Institute for Signal Processing, Technische Universität München (TUM), Germany, in 2021. His research interests include signal processing for millimeter-wave and multiuser communications.



Luis Castedo (Senior Member, IEEE) received the Ph.D. Telecommunications Engineering degree from the Technical University of Madrid, Spain, in 1993. Since 1994, he has been a Faculty Member with the Department of Computer Engineering, University of A Coruña (UDC), Spain, where he became a Professor in 2001 and acted as the Chairman between 2003 and 2009, and again from 2021 to the present. He had previously held several research appointments at the University of Southern California (USC) and École supérieure d'électricité (SUPELEC). His research interests are signal processing, communication theory and prototyping of terminal equipment for wireless communications. He has co-authored more than 300 papers in peer-reviewed international journals and conferences. He has also been the Principal Researcher of more than 50 research projects funded by public organisms and private companies. Between 2014 and 2018, he has been Manager of the Communications and Electronic Technologies (TEC) program in the State Research Agency of Spain. He has been General Co-Chair of the 8th IEEE Sensor Array and Multichannel Signal Processing Workshop in 2014 and the 27th European Signal Processing Conference in 2019. He is recipient of three best paper awards at international conferences and the research medal "Isidro Parga Pondal" awarded by the Royal Galician Academy of Sciences in 2021.