

The Bandit Whisperer: Communication Learning for Restless Bandits

Yunfan Zhao^{1, 2*}, Tonghan Wang^{1*}, Dheeraj Mysore Nagaraj³, Aparna Taneja³, Milind Tambe^{1, 3}

¹Harvard University,

²GE Healthcare,

³Google Deepmind

Abstract

Applying Reinforcement Learning (RL) to Restless Multi-Arm Bandits (RMABs) offers a promising avenue for addressing allocation problems with resource constraints and temporal dynamics. However, classic RMAB models largely overlook the challenges of (systematic) data errors, a common occurrence in real-world scenarios due to factors like varying data collection protocols and intentional noise for differential privacy. We demonstrate that conventional RL algorithms used to train RMABs can struggle to perform well in such settings. To solve this problem, we propose the first communication learning approach in RMABs, where we study which arms, when involved in communication, are most effective in mitigating the influence of such systematic data errors. In our setup, the arms receive Q-function parameters from similar arms as messages to guide behavioral policies, steering Q-function updates. We learn communication strategies by considering the joint utility of messages across all pairs of arms and using a Q-network architecture that decomposes the joint utility. Both theoretical and empirical evidence validate the effectiveness of our method in significantly improving RMAB performance across diverse problems.

1 Introduction

Restless Multi-Arm Bandits (RMABs) have been successfully applied to a range of multi-agent constrained resource allocation problems, such as healthcare, online advertising, and anti-poaching (Ruiz-Hernández, Pinar-Pérez, and Delgado-Gómez 2020; Lee, Lavieri, and Volk 2019; Hodge and Glazebrook 2015; Mate et al. 2022; Modi, Mary, and Moy 2019; Zhao, Krishnamachari, and Liu 2008; Bagheri and Scaglione 2015; Tripathi and Modiano 2019). In classical RMABs, there are a fixed number of arms, each may be described as an MDP with two actions (active and passive). At each timestep, constrained by the global budget, the RMAB strategically selects a subset of arms to play the active action with the objective of maximizing long-term accumulated rewards. Although finding an exact RMAB solution is PSPACE-hard (Papadimitriou and Tsitsiklis 1999), the whittle index solution is known to be near optimal (Whittle 1988; Weber and Weiss 1990) in this situation. Beyond classical RMABs, important

applications of RMABs involve more complex systems with multiple actions, continuous state spaces, and arms opting in and out of the system. Recent studies have explored a variety of methods to address these challenges (Hawkins 2003; Killian, Perrault, and Tambe 2021; Verloop 2016; Zayas-Caban, Jasin, and Wang 2019; Ghosh et al. 2023), with deep Reinforcement Learning (RL) emerging as a particularly effective approach, achieving state-of-the-art results across various RMAB settings (Zhao et al. 2023; Xiong and Li 2023; Wang et al. 2023; Killian et al. 2022).

Unfortunately, these previous studies consider RMAB models where the reward data can be collected from every arm reliably. In real-world scenarios modeled via RMABs, such as maternal healthcare (Biswas et al. 2021; Verma et al. 2024; Behari et al. 2024) and SIS epidemic model concerning COVID intervention (Yaesoubi and Cohen 2011), each arm represents a different entity, which could be a beneficiary or a geographical region, where data collection protocols are potentially different. Different arms may be affected by noise caused by **systematic errors** that arise due to the following practical reasons. **(1) Variability in data collection and handling protocols.** There may be differences in data collection in different arms (entities), e.g., differences in the frequency and quality of data collected, in worker awareness or standards for data preprocessing, cleaning, and digitization. This variability represents sampling bias (Paulus et al. 2023) and may cause systematic errors (Dubrow 2022) in some arms. For example, in maternal care, specific data collection approaches may cause systematic overestimation of expected delivery dates (Fulcher et al. 2020). **(2) Data integrity and privacy.** For example, when data is gathered through surveys, topical or political bias (Paulus et al. 2023) can introduce systematic errors in some arms due to influence of authorities or impacted populations (e.g., keeping the number of COVID cases low (Paulus et al. 2023)). Additionally, driven by the principle of differential privacy, noise might be intentionally added to the data (Dwork, Roth et al. 2014).

Motivated by these scenarios, in this paper we study the RMAB model where some arms are affected by noise in their observed rewards at some states, with these noisy arms and rewards not known a priori. The challenge of applying existing RL methods to this setting is that systematic errors might create false optima, preventing the exploration of states with true high rewards and leading to suboptimal actions

*These authors contributed equally.

selected by the RMAB. For example, the aforementioned systematic overestimation errors of women’s delivery dates can lead to suboptimal resource allocation, thereby reducing the number of deliveries in healthcare facilities (Fulcher et al. 2020).

We propose the first communication learning approach in RMABs to deal with systematic reward errors, enabling arms to learn to communicate useful information to improve local Q-learning. We formally model the communication learning problem as a Multi-Agent MDP (Boutilier 1996), which is in addition to and distinct from per-arm MDP without considering communication. In this framework, each arm has a binary communication action. Choosing to communicate allows an arm to receive the local Q function parameters from a similar arm. Using these parameters, the receiving arm establishes a behavior policy to collect experience samples for updating its local Q function. This behavior policy may enable exploration and help escape false optima. In this way, the communication reward can be calculated as the difference in RMAB return with and without communication. We then adopt Q-learning for communication learning. Since RMABs optimize a centralized objective involving return from all the arms, the communication reward captures the joint influence of communication of all arms. To solve this problem and learn the individual communication strategy of each arm, we propose to use a decomposed Q-network architecture for communication learning that factorizes the joint communication Q function into local communication Q function.

To demonstrate the effectiveness our method, we rigorously compare the sample complexity required to learn a near-optimal per-arm Q function when learning with and without communication. We show theoretically that there exists RMABs where communication from non-noisy to noisy arms can reduce the Q estimation errors exponentially with respect to the number of states. Interestingly, further studies reveal that communication from noisy to non-noisy arms can also be beneficial, provided that the behavior policy, on the receiver arms’ MDP, achieves reasonable coverage over the state-action space and does not decorrelate too slowly from the initial states.

Empirically, we validate our method on synthetic problems, as well as the ARMMAN maternal healthcare (Verma et al. 2024) and the SIS epidemic intervention problem (Yae-soubi and Cohen 2011) built upon real-world data. Results show that our method significantly outperforms both the non-communicative learning approach and the approach with fixed communication strategies, achieving performance levels comparable to those learning in noise-free environments. Results are robust across problems, resources budgets, and noise levels. Furthermore, visualization of the learned communication strategies supports our theoretical findings that communication from noisy arms to non-noisy arms, and vice versa, can be helpful.

Related works. To the best of our knowledge, this is the first work that introduces communication into RMABs. This is perhaps surprising given that communication is an important topic that has been extensively studied in other multi-agent problems. In Dec-POMDPs (Oliehoek, Amato et al.

2016; Kwak et al. 2010; Nair et al. 2004), previous works explore decentralized communication learning (Sukhbaatar, Ferguson et al. 2016; Hoshen 2017; Jiang and Lu 2018; Singh, Jain, and Sukhbaatar 2019; Wang et al. 2019b), the emergence of natural language (Foerster et al. 2016; Das et al. 2017; Lazaridou, Peysakhovich, and Baroni 2017; Mordatch and Abbeel 2018; Zhao et al. 2024b; Kang, Wang, and de Melo 2020), and implicit communication by coordination graphs (Böhmer, Kurin, and Whiteson 2020; Wang et al. 2021; Kang et al. 2022; Yang et al. 2022), influence on each other (Jaques et al. 2019; Wang et al. 2019a), and attention mechanisms (Huang, Mordatch, and Pathak 2020; Kurin et al. 2020; Dong et al. 2022, 2023). Different from these works, we consider communication under noisy returns and in the RMAB setting where the central planner policy constrained by resource budgets determines the influence of communication.

In the literature of cooperative multi-agent reinforcement learning (Yu et al. 2022; Wen et al. 2022; Kuba et al. 2021; Wang et al. 2020a; Christianos, Schäfer, and Albrecht 2020; Peng et al. 2021; Wang et al. 2020c; Jiang et al. 2019; Wang et al. 2020b; Nagaraj et al. 2023), value decomposition methods (Sunehag et al. 2018; Rashid et al. 2018; Son et al. 2019; Li et al. 2021) have recently witnessed great success in addressing challenges such as partial observability and learning scalability. These methods decompose the joint Q-value conditioned on true states and joint actions into per agent local Q functions based on local action-observation history and local actions. In this work, we adopt a linear value decomposition framework (Sunehag et al. 2018) for the communication Q network to solve the unique problems of learning with the presence of noisy arms.

2 Preliminaries

2.1 RMABs with Noisy Arms

We study multi-action Restless Multi-Arm Bandits (RMABs) with N arms. Each arm $i \in [N]$ is associated with a Markov Decision Process $\mathcal{M}_i = (\mathcal{S}_i, \mathcal{A}_i, \mathcal{C}_i, T_i, R_i, \beta, \mathbf{z}_i)$. Here, \mathcal{S}_i is the state space that can be either discrete or continuous. The reward at a state is given by a function $R_i : \mathcal{S}_i \rightarrow \mathbb{R}$, and $\beta \in [0, 1)$ is a discount factor for future rewards. We assume that there is a set of **noisy arms** $\mathcal{N}_\epsilon \subset [N]$, and the rewards of a noisy arm $i \in \mathcal{N}_\epsilon$ are unreliable and affected by noise $\epsilon_i(s_i)$ at some states $s_i \in \mathcal{S}_{\epsilon,i} \subset \mathcal{S}_i$. We consider systematic errors: $\tilde{R}_i(s_i) = R_i(s_i) + \epsilon_i(s_i)$, where $\epsilon_i(s_i)$ is sampled from a distribution \mathcal{E}_i and fixed during training and testing. The identity of noisy arms \mathcal{N}_ϵ , the affected states $\mathcal{S}_{\epsilon,i}$, the reward noise ϵ_i , and its distribution \mathcal{E}_i are not known a priori.

\mathcal{A}_i is a finite set of discrete actions. Each action $a \in \mathcal{A}_i$ has a cost $\mathcal{C}_i(a)$. Following the standard bandit assumption that an arm can be “not pulled” at no cost, we set $\mathcal{C}_i(0) = 0$. The probability of transitioning from one state to another given an action is specified by the function $T_i : \mathcal{S}_i \times \mathcal{A}_i \times \mathcal{S}_i \rightarrow [0, 1]$. In line with previous work (Zhao et al. 2023; Elmachtoub, Gupta, and Zhao 2023; Elmachtoub et al. 2023; Zhao et al. 2022) we assume that $\mathcal{S}_i, \mathcal{A}_i$, and \mathcal{C}_i are the same for all arms $i \in [N]$ and omit the subscript i for simplicity. Each arm has a feature vector $\mathbf{z}_i \in \mathbb{R}^m$ that provides useful information about the transition dynamics of arm i . We define

the arm with the most similar features with arm i as $\nu_i = \arg \min_j \|z_j - z_i\|_2$.

We consider a constrained global resource setting where, at every timestep $t \in [H]$, where H is the horizon, the total cost of actions taken is no greater than a given budget B . Under this constraint, the *central planner* selects one action for each of the N arms at each timestep. Formally, the central planner has a policy π to maximize its expected discounted return $G(\pi) = \mathbb{E}_\pi[\sum_{t=1}^H \beta^t \sum_{i \in [N]} R_i(s_i^t)]$, where s_i^t is the state of arm i at time step t . π is typically computed by solving the constrained Bellman equation (Hawkins 2003; Killian, Perrault, and Tambe 2021; Killian et al. 2022):

$$\begin{aligned} J(\mathbf{s}) &= \sum_{i=1}^N R_i(s_i) + \beta \max_{\mathbf{a}} \mathbb{E}_{s'} [J(s') \mid \mathbf{s}, \mathbf{a}] \\ \text{s.t. } \sum_{i=1}^N \mathcal{C}(a_i) &\leq B. \end{aligned} \quad (1)$$

2.2 Reinforcement Learning in RMABs

Deep reinforcement learning has been proven to be efficient in solving the RMAB problem (Avrachenkov and Borkar 2022; Killian et al. 2022; Zhao et al. 2024a). To learn arm policies, we take the Lagrangian relaxation of the budget constraint (Hawkins 2003), defining

$$J(\mathbf{s}, \lambda) = \frac{\lambda B}{1 - \beta} + \sum_{i=1}^N \max_{j \in [A]} \{Q_i(s_i, j)\}, \quad (2)$$

where Q_i is arm i 's Q function estimating its expected discount return of an action at a state. We use deep Q-learning and parameterize the action-value function $Q_i(s_i, a_i; \theta_i)$ of each arm i with θ_i , which is updated by the TD-loss (Sutton and Barto 2018; Killian, Perrault, and Tambe 2021):

$$\begin{aligned} \mathcal{L}_{\text{TD}}(\theta_i) &= \mathbb{E}_{(s_i, a_i, s'_i) \sim \mathcal{D}} \left[\left(R_i(s_i) - \lambda \mathcal{C}(a_i) \right. \right. \\ &\quad \left. \left. + \beta \max_a Q_i(s'_i, a; \theta_i) - Q_i(s_i, a_i; \tilde{\theta}_i) \right)^2 \right], \end{aligned} \quad (3)$$

where $Q_i(\cdot; \tilde{\theta}_i)$ is a target network whose parameters are periodically copied from Q_i . The expectation means that the samples come from a replay buffer \mathcal{D} .

To *infer* a central planner policy that adheres to the budget constraint, previous research (Killian et al. 2022) constructs a knapsack problem using $Q_i(s_i, a_i)$ as values and $\mathcal{C}(a_i)$ as weights, within the constraints defined in the optimization problem 1. The solution to this knapsack problem identifies the joint action that maximizes the sum of the learned Q_i values while conforming to the budget constraint. In this case, we denote the planner policy by $\pi(\theta)$ to highlight its dependency on the joint parameters of individual Q-functions $\theta = (\theta_1, \dots, \theta_N)$.

3 RMAB Learning with Communication

In this section, we introduce our method of learning communication to share useful knowledge among arms. We first define the communication problem as a Multi-Agent

MDP (Boutilier 1996) and address the following problems: (1) What are the messages? (Sec. 3.1) How do they affect local Q learning? (Sec. 3.2) (2) How to generate reward signals that measure the influence of communication? (Sec. 3.3) (3) How to use the communication reward to learn communication strategies? (Sec. 3.4) (4) Why and under what conditions can communication help learning? (Sec. 3.2 and 4)

3.1 Define the Communication Learning Problem

Akin to traditional deep RMAB learning algorithms (Fu et al. 2019; Biswas et al. 2021; Killian et al. 2022; Avrachenkov and Borkar 2022), our approach involves each arm learning its local Q networks $Q_i(s_i, a_i; \theta_i)$ to predict the expected return of taking action a_i in state s_i . The central planner then selects centralized actions based on the local Q values.

When the data source is noisy, the local Q functions may be inaccurate, leading to the central planner choosing sub-optimal actions. Figure 3 in Appendix C.3 shows such an example. To address this issue, we introduce a mechanism that allows each arm to learn a communication policy that diminishes the impact of noisy data and enhances the overall effectiveness of the central planner.

We define the communication learning problem within the framework of a Multi-Agent MDP (Boutilier 1996).

Definition 1 (The Communication MDP). *The communication MDP in the RMAB setting is modelled by a multi-agent MDP $\mathcal{M}^c = (\mathcal{S}^c, I, \mathcal{A}^c, T^c, r^c)$. I is the set of players, which in our case is the set of arms. $\mathcal{S}^c = \Theta$ is the state space, where Θ is the joint parameter space of Q networks. $\theta = (\theta_1, \dots, \theta_N) \in \mathcal{S}^c$ corresponds to joint Q function parameters. $\mathcal{A}^c = \{0, 1\}^J$ is the action space. $a_i^c = 1$ means arm i decides to ask Q network parameters from the arm with the most similar features, ν_i , as the message. If $a_i^c = 0$, arm i will not receive a message. The joint communication action \mathbf{a}^c leads to a transition in θ (Eq.6) and a reward r^c (Eq. 7).*

The definition of this Communication MDP is associated with the underlying RMAB in the sense that the Communication MDP solution affects the action selection of the central planner. After every K learning epochs of Q_i in the MDP \mathcal{M}_i , parameters of arms' Q functions are used as the state for this Communication MDP \mathcal{M}^c . We assume that each arm i in the system has the chance to communicate with the arm with the most similar features, denoted by arm ν_i . This results in N possible communication channels (i.e. arm i will receive a message if $a_i^c = 1$). In this way, the set of players in \mathcal{M}^c is the set of arms $[N]$. In addition to selecting actions in the arm MDP \mathcal{M}_i , each arm also selects actions in this Communication MDP \mathcal{M}^c . These *communication actions* determine the communication channel activation pattern, which remains frozen for the next K learning epochs. We will discuss the details in the following subsections.

An alternative design option is to adopt a denser communication scheme, i.e., an arm can receive messages from multiple other arms. As formally detailed in Proposition 4 in Sec. 4, under certain assumptions regarding concentratability (Munos 2003; Chen and Jiang 2019), our sparse communication scheme is shown to effectively learn near-optimal value functions with the same errors as the dense scheme.

Therefore, we favor the sparse scheme due to its lower demands on communication bandwidth.

3.2 Communication Actions and Transition Dynamics

We now describe how the communication action \mathbf{a}^c induces a state transition in the space \mathcal{S}^c . Specifically, when a_i^c is 1, arm i receives a message from arm ν_i , which is the parameters of the local Q network of arm ν_i , θ_{ν_i} . Using this message, arm i formulates a behavior policy π_{ν_i} as follows:

$$\pi_{\nu_i}(\cdot|s_i) = \text{SoftMax}(Q_{\nu_i}(s_i, \cdot)). \quad (4)$$

Using behavior policy π_{ν_i} , arm i navigates in its arm MDP \mathcal{M}_i , collecting transition samples in a replay buffer $\mathcal{D}_{\pi_{\nu_i}}$. This buffer is then used to update the local Q-network. Therefore, the TD-loss for arm i with the communication channel i activated is given by:

$$\begin{aligned} \mathcal{L}_{\text{TD}}^c(\theta_i|a_i^c=1) = & \mathbb{E}_{(s_i, a_i, s'_i) \sim \mathcal{D}_{\pi_{\nu_i}}} \left[\left(R_i(s_i) - \lambda c_{a_i} \right. \right. \\ & \left. \left. + \beta \max_a Q_i(s'_i, a; \theta_i) - Q_i(s_i, a_i; \tilde{\theta}_i) \right)^2 \right]. \end{aligned}$$

Alternatively, if arm i opts not to receive communication, it engages in the ϵ -greedy strategy to explore the MDP \mathcal{M}_i and generate the replay buffer. Therefore,

$$\mathcal{L}_{\text{TD}}^c(\theta_i|a_i^c=0) = \mathcal{L}_{\text{TD}}(\theta_i). \quad (5)$$

As such, the updates to local Q parameters are dependent on the communication action:

$$\theta'_i(a_i^c) = \theta_i - \alpha \nabla_{\theta_i} \mathcal{L}_{\text{TD}}^c(\theta_i|a_i^c), \quad (6)$$

where α is the learning rate. Eq. 6 specifies the state transition dynamics of the Communication MDP.

3.3 Communication Reward

Lemma 5 and 6 in Sec. 4 reveal that some seemingly reasonable fixed communication patterns, like those from a non-noisy arm to a noisy arm, can lead to an exponential higher sample complexity for learning a near-optimal Q function in the noisy arm. In fact, as shown in Proposition 2 of Sec. 4, the influence of communication is determined by the state-action space coverage and the mixing time of the behavior policy π_{ν_i} on the receiver arms.

These findings highlight the intricate and dynamic nature of the communication learning problem in RMAB. In practice, to adaptively decide the optimal communication action, we propose that each arm learns a communication Q function Q_i^c . This approach encounters obstacles in the RMAB setting, where clear, separate reward signals for each Q_i^c are absent. The central planner policy instead bases action decisions on the updated Q-functions of all arms and the budget. The collaborative contribution of the joint communication action \mathbf{a}^c to the learning objective is:

$$r^c(\theta, \mathbf{a}^c) = G_T(\pi(\theta'(\mathbf{a}^c))) - G_T(\pi(\theta'(\mathbf{0}))). \quad (7)$$

Here, $G_T(\pi) = \frac{1}{T} \sum_{k=1}^T [\sum_{t=1}^H \beta^t \sum_{i \in [N]} R_i(s_i^{k,t})]$ is the empirical return under policy π averaged over T episodes,

where $s_i^{k,t}$ is arm i state at timestep t of episode k . $\pi(\theta'(\mathbf{a}^c))$ and $\pi(\theta'(\mathbf{0}))$ are the planner policies with the arm Q-functions updated with and without messages, respectively. The communication reward is defined as the difference in the planner's return in these two cases.

3.4 Decomposed Communication Q Function

With the definition of the communication reward (Eq. 7), the simplest strategy of communication learning is to estimate a *joint* communication action-value function $Q^c(\theta, \mathbf{a}^c)$. However, the question is that there are N communication channels, and even though our communication action is binary, there are $O(2^N)$ possible joint actions \mathbf{a}^c . With this exponential number of actions, it is challenging to conduct deep Q-learning, which requires a maximization operation over the whole action space.

We propose to solve this problem by value decomposition (Sunehag et al. 2018). Specifically, we represent the global communication Q-function as a summation of the channel-level, individual communication Q-functions:

$$Q^c(\theta, \mathbf{a}^c; \xi) = \sum_{i \in [N]} Q_i^c(\theta, a_i^c; \xi_i). \quad (8)$$

Here, ξ_i is the parameters of the channel-level communication Q-function, and ξ is the joint parameter consisting of all ξ_i . Such a decomposition satisfies an important property:

$$\arg \max_{\mathbf{a}^c} Q^c = \left(\arg \max_{a_1^c} Q_1^c, \dots, \arg \max_{a_N^c} Q_N^c \right), \quad (9)$$

which is because Eq. 8 allows $\partial Q^c / \partial Q_i^c \geq 0$. This means that the joint action that maximizes the Q^c can be obtained by independently selecting the local communication action that maximize Q_i . As a whole, network $Q^c(\cdot; \xi)$ is a standard Q-network conditioned on states θ and actions \mathbf{a}^c , allowing it to be trained using joint rewards r^c (Eq. 7). On a finer level, it has a decomposable structure simplifying the $\arg \max$ calculation. Therefore, $Q^c(\cdot; \xi)$ can be trained by the following TD-loss:

$$\begin{aligned} \mathcal{L}_{\text{TD}}^d(\xi) = & \mathbb{E}_{(\theta, \mathbf{a}^c, \theta') \sim \mathcal{D}} \left[\left(r^c(\theta, \mathbf{a}^c) \right. \right. \\ & \left. \left. + \beta \sum_{i \in [N]} \max_a Q_i^c(\theta', a; \tilde{\xi}_i) - \sum_{i \in [N]} Q_i^c(\theta, a_i^c; \xi_i) \right)^2 \right]. \end{aligned}$$

$Q_i^c(\cdot; \tilde{\xi}_i)$ is a target network with parameters periodically copied from Q_i^c , and \mathcal{D} is a replay buffer. In this way, we implicitly achieve credit assignment, because we use a joint reward r^c for training but get a set of Q functions for individual arms, which can recover the best joint action locally.

4 Theoretical Analysis

We begin by investigating the influence of communication and establishing the specific conditions under which communication can reduce the sample complexity for the receiver arm in learning a near-optimal Q function. We define the minimum state-action occupancy probability,

$$\mu_{\min} := \min_{s,a} \mu_{\pi_{\nu_i}}(s, a), \quad (10)$$

as the lowest probability of any state-action pair under the stationary distribution $\mu_{\pi_{\nu_i}}$ induced by the behavior policy π_{ν_i} on the receiver MDP, where ν_i is the arm with the most similar feature to arm i . We further define the mixing time (Hsu, Kontorovich, and Szepesvári 2015; Li et al. 2020) $t_{\text{mix}} :=$

$$\min \left\{ t \mid \max_{s_0, a_0} \sup_{s, a} |P^t(s, a | s_0, a_0) - \mu_{\pi_{\nu_i}}(s, a)| \leq \frac{1}{4} \right\}, \quad (11)$$

where P^t is the state-action distribution at timestep t given the initial state and action. The mixing time t_{mix} captures how fast the behavior policy decorrelates from the initial state. We now show the condition under which communication can be helpful.

Proposition 2 (Useful Communication). *For the Q-learning of arm i , when*

$$\mu_{\text{min}} > \frac{2(1-\beta)^2}{|\mathcal{S}||\mathcal{A}|}, \text{ and } t_{\text{mix}} \leq \frac{1}{\epsilon_e^2(1-\beta)^4}, \quad (12)$$

with a probability at least $(1-\delta)^2$ for any $0 < \delta < 1$, compared to learning without communication, choosing to receive communication and using behavior policy π_{ν_i} leads to better sample complexity in the worst case for learning ϵ_e -optimal Q function such that $\|Q_i(s, a) - Q_i^*(s, a)\| < \epsilon_e, \forall (s, a)$.

The Proof is detailed in Appendix C.2 Proposition 2 shows that communication can accelerate arm i 's Q-learning as long as π_{ν_i} visits every state-action pair with a small probability and is not too slowly disentangling from the initial state. This finding is solely dependent upon the characteristics of the behavior policy in the receiver MDP \mathcal{M}_i and does not require the sender or the receiver to be non-noisy or noisy.

We then analyze the effectiveness of a sparse communication strategy. Specifically, we compare our setup where each arm i receives messages from an arm ν_i against a dense communication scheme where arm i receives messages from d other arms. In the dense communication scheme, the behavior policy $\pi_{\bar{\nu}_i}$ is the average of policies from all senders. Specifically, we collect the same number of samples using each sender's policy and amalgamate them into a unified experience replay buffer.

The following assumption on concentratability asserts that both π_{ν_i} and $\pi_{\bar{\nu}_i}$ provide a reasonable coverage over the state-action space.

Assumption 3 (Concentratability Coefficient). *There exists $C_{\nu_i} < \infty, C_{\bar{\nu}_i} < \infty$ s.t.*

$$\frac{v(s, a)}{\mu_{\bar{\nu}_i}(s, a)} \leq C_{\bar{\nu}_i} \text{ and } \frac{v(s, a)}{\mu_{\nu_i}(s, a)} \leq C_{\nu_i}, \text{ for } \forall (s, a) \in (\mathcal{S}, \mathcal{A}),$$

for any admissible distribution v that can be generated by following some policy for some timesteps in the MDP \mathcal{M}_i . Here μ_{ν_i} and $\mu_{\bar{\nu}_i}$ are the stationary distribution of π_{ν_i} and $\pi_{\bar{\nu}_i}$ on \mathcal{M}_i , respectively.

Under this assumption, we can prove that the sparse communication scheme can also learn a near-optimal Q function as under the dense scheme. The key observation that underpins the proof is $\frac{1}{d}\mu_{\nu_i}(s, a) \leq \mu_{\bar{\nu}_i}(s, a)$, where $d < N$ is the

maximum number of other arms that can send messages to arm i . The convergence of the communication learning under the dense scheme implies that $C_{\bar{\nu}_i}$ is finite, which further indicates the finiteness of C_{ν_i} and thus the convergence of the communication learning of the sparse scheme.

Proposition 4 (Sparse Communication). *For a dense communication strategy, there exists a sparse communication strategy which approximates the optimization step in Equation (6) faithfully. Additionally, both these two strategies can learn near-optimal value functions, and the sparse communication scheme requires at most $O(d)$ times more samples to achieve this.*

Detailed proof is deferred to Appendix C.1. We next show that a communication channel from a non-noisy arm to a noisy arm can be either helpful (Lemma 5) or detrimental (Lemma 6), depending on the specific structures of the sender and receiver MDPs.

Lemma 5. *There exists a family of RMABs where, without communication from non-noisy arms to noisy arms, the error in Q values of a noisy arm could be $\max_{s, a} |Q_i^*(s, a) - Q_i(s, a)| = O\left(\left(\frac{1}{\beta}\right)^{|\mathcal{S}|}\right)$ even after $O\left(\left(\frac{1}{\epsilon}\right)^{|\mathcal{S}|}\right)$ epochs. $\epsilon \in (0, 1)$ is the exploration rate in ϵ -greedy exploration scheme. By contrast, this Q error can be reduced to 0 with communication from non-noisy arms to noisy arms.*

Lemma 6. *There exists a family of RMABs where, with communication from a non-noisy arm to a noisy arm, we need $O\left(\left(\frac{1}{\epsilon}\right)^{|\mathcal{S}|}\right)$ times more samples to learn a near-optimal Q function in the noisy arm than without communication.*

The proofs of Lemma 5 and 6 can be found in Appendix C.3. The observation that the errors in the Q function and the sample complexity of Q learning can increase exponentially with the number of states in the worst case scenarios highlights the need to move away from predefined communication channels. Instead, communication Q-learning provides a more adaptable and effective approach.

5 Experiments

We present experimental evaluations and ablation studies to investigate: (1) The learning performance of our communication method compared to baselines; (2) The characteristics of learnt communication strategies; (3) The robustness of our method in terms of resource budgets, the number of arms, noise types and levels. We discuss additional experimental details and hyperparameters in Appendix B.

We evaluate our methods on the following domains. In all domains, 80% of the arms are noisy, and noisy arms observe reward drawn from a perturbed reward function $\tilde{R}(s) = R(s) + \epsilon(s)$, where $\epsilon \sim \mathcal{N}(0, 1)$. In ablation studies (Tables 2 and 3), we test various numbers of noisy arms and noise type.

Synthetic: We consider a synthetic RMAB with continuous states. For state $s_i \in [0, 1]$ and action $a \in \{0, 1\}$ of arm i , the next state s'_i is determined by: $s'_i = \text{clip}(s_i + \mathcal{N}(\mu_{ia}, \sigma_{ia}), 0, 1)$. Here transition parameters are sampled uniformly, specifically $\mu_{i1} = -\mu_{i0} \in [-0.2, 0.2]$, $\sigma_{i1} = \sigma_{i0} \in [0.1, 0.2]$.

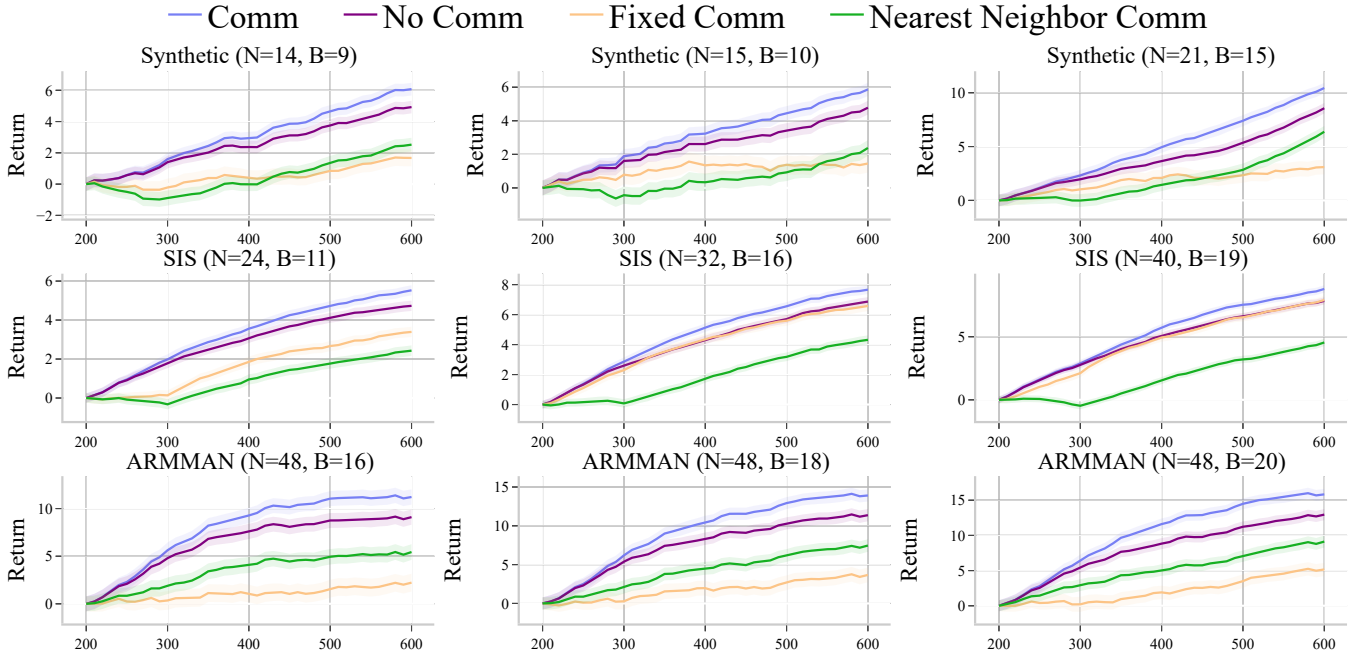


Figure 1: Performance (interquartile mean (Agarwal et al. 2021) and standard error of return over 200 random seeds) of our method, baselines, and ablations in three environments with different numbers of arms N and resource budgets B . Communication learning starts at epoch 200. We present results on **two additional baseline** in Tables 1 and 2 (see Appendix B for more ablations).

SIS Epidemic Model (Yaesoubi and Cohen 2011; Killian et al. 2022; Zhao et al. 2024a): Each arm p represents a sub-population in a geographic region, with the number of uninfected individuals being the state s_p . We implement a budget-constrained intervention framework with three actions: a_0 represents no intervention, a_1 requires weekly COVID testing, and a_2 involves distributing face masks. Actions change the disease transmission parameters in a targeted manner, reflecting their real-world public-health impacts. The reward noise is due to variability in data collection, which may be caused by poor COVID testing equipment (Embrett et al. 2022; Schmitt-Grohé, Teoh, and Uribe 2020) and influence of authorities or impacted populations (keeping the COVID numbers low) (Paulus et al. 2023).

ARMMAN: The dataset is collected by ARMMAN (ARMMAN 2019), an NGO in India aimed at improving health awareness for a million expectant and new mothers via automated voice messaging programs. Each arm represents a cluster of mothers in a geographical location. Actions are sending a health worker to visit mothers in a location to improve engagement. States, continuous valued in $[0,1]$, represent the average engagement of mothers in a location (Danassis et al. 2023). The transition dynamics is established empirically from the data (data usage and consent are in Appendix A.2). Systematic reward errors are present due to difference in infrastructure such as cellular network and electricity (Kumar et al. 2022) (e.g. a geographic location may have poor cellular signals or frequent power outage, causing a systematic underestimation of mothers’ engagement).

Learning Performance. Figure 1 shows the experimen-

tal results in Synthetic, SIS, and ARMMAN where we compare our method Comm (Ours) against various baselines and ablations. No Comm is an ablation where the RMAB with noisy arms is trained without communication. Fixed Comm involves a preset, fixed communication pattern where at each timestep, each noisy arm receives information from a similar (ℓ_2 distance of features), noise-free arm. Fixed Comm uses oracle information about which arms are noisy, which is not known by other methods. Nearest Neighbor Comm is a baseline strategy where an arm always receives communication from the nearest neighbor in feature space, regardless of whether the sender/receiver is a noisy arm. We include **two additional baselines** in Tables 1 and 2.

Our method significantly outperforms No Comm, Fixed Comm, and Nearest neighbor Comm. Notably, the Fixed Comm strategy, despite the oracle information, often performs worse than No Comm. This observation aligns with our theoretical analysis in Proposition 2 showing that communication channels from non-noisy to noisy arms are not always helpful. Our method consistently outperforms across various number of arms, budgets, and domains. For instance, in the Synthetic environment with $N=21$ and $B=15$, our method achieves a return of approximately 10 at epoch 600, compared to about 8 for the no communication (No Comm) baseline, representing an 20 percentage point improvement. Similarly, in the ARMMAN environment with $N=48$ and $B=20$, our method reaches a return of 15, outperforming the No Comm baseline’s return of approximately 12.5, which is a 20 percentage point increase. These results demonstrate the effectiveness of our

Method	Synthetic (N=21, b=15)			SIS (N=32, B=16)			ARMMAN (N=48, B=20)		
	100%	110%	120%	100%	110%	120%	100%	110%	120%
No Noise	100%	100%	100%	100%	100%	100%	100%	100%	100%
Comm	75%	53%	38%	63%	55%	51%	77%	72%	68%
No Comm	62%	39%	24%	55%	45%	42%	63%	59%	52%
Fixed/Comm	21%	6%	12%	56%	42%	46%	24%	15%	27%
Nearest Neighbor Comm	44%	20%	3%	32%	22%	23%	46%	43%	31%
Random Comm	-24%	-45%	-58%	22%	13%	9%	-22%	-26%	-35%

Table 1: Average returns (100 seeds) with increased noise, normalized with respect to No Noise and Comm (ours) before communication (see Remark 7 in Appendix B). We test Gaussian Noise $\mathcal{N}(0, \sigma)$. $x\%$ noise level means $\sigma = x/100$.

Method	Synthetic (N=21, B=15)			SIS (N=32, B=16)			ARMMAN (N=48, B=20)		
	15	16	17	22	24	26	34	36	38
No Noise	100%	100%	100%	100%	100%	100%	100%	100%	100%
Comm	78%	78%	75%	69%	71%	63%	82%	80%	77%
No Comm	67%	66%	62%	61%	61%	55%	68%	59%	63%
Fixed/Comm	63%	45%	21%	65%	62%	56%	43%	49%	24%
Nearest Neighbor Comm	42%	47%	44%	11%	29%	32%	47%	37%	46%
Random Comm	-46%	-32%	-24%	-8%	17%	22%	-60%	-30%	-22%

Table 2: Performance (return averaged across 100 random seeds) with numbers of noisy arms, normalized with respect to No Noise and Comm (ours) before communication starts (see Remark 7 in Appendix B).

communication algorithm in handling systematic errors and enhancing learning in noisy RMAB environments.




Figure 2: The change in the proportion of noise-free senders and receiver within all activated channels during the learning process. Results are for Synthetic (N=21, B=15), SIS (N=32, B=16), and ARMMAN (N=48, B=16).

Communication Pattern. Figure 2 illustrates how our learnt communication strategy evolves. We make several observations. First, as the number of learning epochs increase, we observe more information are sent from non-noisy arms to noisy arms. During the early stages of learning, when arms are primarily exploring, communication from noisy arms to non-noisy arms can be beneficial, as shown in Proposition 2. As the learning progresses, the information from noisy sources may become less helpful as the arms have explored sufficiently. As a result, the proportion of communication from non-noisy arms increases. This observation is consistent across the three experimental environments.

Second, notice the percentage of noise-free receivers do not drop to zero and noise-free senders are not at 100% even at the end of training, indicating value of communication from noisy arms. Finally, the differing rates of noise-free arms as senders and receivers across domains indicate that our communication approach adapts to the domain.

Robustness. Tables 1 and 2 show the performance of different methods under various settings (see Appendix B for more ablations). In No Noise, all arms receive accurate, true reward signals without any noise. Arms learn individually on their own data without any communication. In Random Comm, each arm receives communication from a randomly chosen arm. The results demonstrate that our method is robust under increasing noise levels, numbers of noisy arms, and different noise distributions. In contrast, the performance of baselines No Comm and Fixed Comm changes significantly when the environment changes. In particular, on Synthetic, when the number of noisy arms increases from 15 to 17 (Table 2), our Comm method has **only 3%** drop in performance, while Fixed Comm has 42 percentage points drop in performance, despite that Fixed Comm uses ground truth information about which arms are noisy that is not known to our Comm method. Similar trends can be seen on SIS and ARMMAN. These results demonstrate that, compared to a fixed communication protocol with access to oracle information, our approach for communication learning provides better robustness. Additionally, the performance of No Comm supports the argument that conventional RL methods struggle in noisy RMABs.

6 Closing Remarks

The study of RMABs in environments with noisy and unreliable data sources underscores the need for innovative solutions in RL methods. We introduce a novel communication learning algorithm to alleviate the influence of noisy data so that the central planner can make better coordination decisions. Theoretically, we prove that communication can help reduce errors in value function approximation if the behavioral policy’s state-action occupancy probability is lower bounded by a suitable constant. Empirically, we show that our approach significantly improves RMAB performance.

Acknowledgments

The work is supported by Harvard HDSI funding

References

- Agarwal, R.; Schwarzer, M.; Castro, P. S.; Courville, A. C.; and Bellemare, M. 2021. Deep reinforcement learning at the edge of the statistical precipice. *Advances in neural information processing systems*, 34: 29304–29320.
- ARMMAN. 2019. Assessing the Impact of Mobile-based Intervention on Health Literacy among Pregnant Women in Urban India. <https://armman.org/wp-content/uploads/2019/09/Sion-Study-Abstract.pdf>. Accessed: 2022-08-12.
- Avrachenkov, K. E.; and Borkar, V. S. 2022. Whittle index based Q-learning for restless bandits with average reward. *Automatica*, 139: 110186.
- Bagheri, S.; and Scaglione, A. 2015. The restless multi-armed bandit formulation of the cognitive compressive sensing problem. *IEEE Transactions on Signal Processing*, 63(5): 1183–1198.
- Behari, N.; Zhang, E.; Zhao, Y.; Taneja, A.; Nagaraj, D.; and Tambe, M. 2024. A Decision-Language Model (DLM) for Dynamic Restless Multi-Armed Bandit Tasks in Public Health. *arXiv preprint arXiv:2402.14807*.
- Biswas, A.; Aggarwal, G.; Varakantham, P.; and Tambe, M. 2021. Learn to intervene: An adaptive learning policy for restless bandits in application to preventive healthcare. *arXiv preprint arXiv:2105.07965*.
- Böhmer, W.; Kurin, V.; and Whiteson, S. 2020. Deep coordination graphs. In *International Conference on Machine Learning*, 980–991. PMLR.
- Boutilier, C. 1996. Planning, learning and coordination in multiagent decision processes. In *TARK*, volume 96, 195–210. Citeseer.
- Chen, J.; and Jiang, N. 2019. Information-theoretic considerations in batch reinforcement learning. In *International Conference on Machine Learning*, 1042–1051. PMLR.
- Christianos, F.; Schäfer, L.; and Albrecht, S. 2020. Shared experience actor-critic for multi-agent reinforcement learning. *Advances in neural information processing systems*, 33: 10707–10717.
- Danassis, P.; Verma, S.; Killian, J. A.; Taneja, A.; and Tambe, M. 2023. Limited resource allocation in a non-Markovian world: the case of maternal and child healthcare. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, 5950–5958.
- Das, A.; Kottur, S.; Moura, J. M.; Lee, S.; and Batra, D. 2017. Learning cooperative visual dialog agents with deep reinforcement learning. In *Proceedings of the IEEE International Conference on Computer Vision*, 2951–2960.
- Dong, H.; Wang, T.; Liu, J.; and Zhang, C. 2022. Low-rank modular reinforcement learning via muscle synergy. *Advances in Neural Information Processing Systems*, 35: 19861–19873.
- Dong, H.; Zhang, J.; Wang, T.; and Zhang, C. 2023. Symmetry-aware robot design with structured subgroups. In *International Conference on Machine Learning*, 8334–8355. PMLR.
- Dubrow, J. K. 2022. Local data and upstream reporting as sources of error in the administrative data undercount of Covid 19. *International Journal of Social Research Methodology*, 25(4): 471–476.
- Dwork, C.; Roth, A.; et al. 2014. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4): 211–407.
- Elmachtoub, A.; Gupta, V.; and Zhao, Y. 2023. Balanced off-policy evaluation for personalized pricing. In *International Conference on Artificial Intelligence and Statistics*, 10901–10917. PMLR.
- Elmachtoub, A. N.; Lam, H.; Zhang, H.; and Zhao, Y. 2023. Estimate-then-optimize versus integrated-estimation optimization: A stochastic dominance perspective. *arXiv preprint arXiv:2304.06833*.
- Embrett, M.; Sim, S. M.; Caldwell, H. A.; Boulos, L.; Yu, Z.; Agarwal, G.; Cooper, R.; Aj, A. J. G.; Bielska, I. A.; Chishtie, J.; et al. 2022. Barriers to and strategies to address COVID-19 testing hesitancy: a rapid scoping review. *BMC public health*, 22(1): 750.
- Foerster, J.; Assael, I. A.; de Freitas, N.; and Whiteson, S. 2016. Learning to communicate with deep multi-agent reinforcement learning. In *Advances in Neural Information Processing Systems*, 2137–2145.
- Fu, J.; Nazarathy, Y.; Moka, S.; and Taylor, P. G. 2019. Towards q-learning the whittle index for restless bandits. In *2019 Australian & New Zealand Control Conference (ANZCC)*, 249–254. IEEE.
- Fulcher, I.; Hedt, K.; Marealle, S.; Tibaijuka, J.; Abdalla, O.; Hofmann, R.; Layer, E.; Mitchell, M.; and Hedt-Gauthier, B. 2020. Errors in estimated gestational ages reduce the likelihood of health facility deliveries: results from an observational cohort study in Zanzibar. *BMC Health Services Research*, 20: 1–10.
- Ghosh, A.; Nagaraj, D.; Jain, M.; and Tambe, M. 2023. Indexability is Not Enough for Whittle: Improved, Near-Optimal Algorithms for Restless Bandits. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, 1294–1302.
- Hawkins, J. T. 2003. *A Lagrangian decomposition approach to weakly coupled dynamic optimization problems and its applications*. Ph.D. thesis, Massachusetts Institute of Technology.
- Hodge, D. J.; and Glazebrook, K. D. 2015. On the asymptotic optimality of greedy index heuristics for multi-action restless bandits. *Advances in Applied Probability*, 47(3): 652–667.
- Hoshen, Y. 2017. Vain: Attentional multi-agent predictive modeling. In *Advances in Neural Information Processing Systems*, 2701–2711.
- Hsu, D. J.; Kontorovich, A.; and Szepesvári, C. 2015. Mixing time estimation in reversible markov chains from a single sample path. *Advances in neural information processing systems*, 28.

- Huang, W.; Mordatch, I.; and Pathak, D. 2020. One policy to control them all: Shared modular policies for agent-agnostic control. In *International Conference on Machine Learning*, 4455–4464. PMLR.
- Jaques, N.; Lazaridou, A.; Hughes, E.; Gulcehre, C.; Ortega, P.; Strouse, D.; Leibo, J. Z.; and De Freitas, N. 2019. Social influence as intrinsic motivation for multi-agent deep reinforcement learning. In *International conference on machine learning*, 3040–3049. PMLR.
- Jiang, J.; Dun, C.; Huang, T.; and Lu, Z. 2019. Graph Convolutional Reinforcement Learning. In *International Conference on Learning Representations*.
- Jiang, J.; and Lu, Z. 2018. Learning attentional communication for multi-agent cooperation. In *Advances in Neural Information Processing Systems*, 7254–7264.
- Kang, Y.; Wang, T.; and de Melo, G. 2020. Incorporating pragmatic reasoning communication into emergent language. *Advances in neural information processing systems*, 33: 10348–10359.
- Kang, Y.; Wang, T.; Yang, Q.; Wu, X.; and Zhang, C. 2022. Non-linear coordination graphs. *Advances in Neural Information Processing Systems*, 35: 25655–25666.
- Killian, J. A.; Perrault, A.; and Tambe, M. 2021. Beyond “To Act or Not to Act”: Fast Lagrangian Approaches to General Multi-Action Restless Bandits. In *AAMAS*, 710–718. UK: AAMAS.
- Killian, J. A.; Xu, L.; Biswas, A.; and Tambe, M. 2022. Restless and uncertain: Robust policies for restless bandits via deep multi-agent reinforcement learning. In *Uncertainty in Artificial Intelligence*, 990–1000. PMLR.
- Kuba, J. G.; Chen, R.; Wen, M.; Wen, Y.; Sun, F.; Wang, J.; and Yang, Y. 2021. Trust Region Policy Optimisation in Multi-Agent Reinforcement Learning. In *International Conference on Learning Representations*.
- Kumar, S. K. A.; Ihita, G. V.; Chaudhari, S.; and Arumugam, P. 2022. A survey on rural internet connectivity in India. In *2022 14th International Conference on COMMunication Systems & NETWORKS (COMSNETS)*, 911–916. IEEE.
- Kurin, V.; Igl, M.; Rocktäschel, T.; Boehmer, W.; and Whiteson, S. 2020. My Body is a Cage: the Role of Morphology in Graph-Based Incompatible Control. In *International Conference on Learning Representations*.
- Kwak, J.-y.; Yang, R.; Yin, Z.; Taylor, M. E.; and Tambe, M. 2010. Teamwork and coordination under model uncertainty in DEC-POMDPs. In *Workshops at the Twenty-Fourth AAAI Conference on Artificial Intelligence*.
- Lazaridou, A.; Peysakhovich, A.; and Baroni, M. 2017. Multi-agent cooperation and the emergence of (natural) language. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Lee, E.; Lavieri, M. S.; and Volk, M. 2019. Optimal screening for hepatocellular carcinoma: A restless bandit model. *Manufacturing & Service Operations Management*, 21(1): 198–212.
- Li, C.; Wang, T.; Wu, C.; Zhao, Q.; Yang, J.; and Zhang, C. 2021. Celebrating diversity in shared multi-agent reinforcement learning. *Advances in Neural Information Processing Systems*, 34: 3991–4002.
- Li, G.; Wei, Y.; Chi, Y.; Gu, Y.; and Chen, Y. 2020. Sample complexity of asynchronous Q-learning: Sharper analysis and variance reduction. *Advances in neural information processing systems*, 33: 7031–7043.
- Mate, A.; Madaan, L.; Taneja, A.; Madhiwalla, N.; Verma, S.; Singh, G.; Hegde, A.; Varakantham, P.; and Tambe, M. 2022. Field Study in Deploying Restless Multi-Armed Bandits: Assisting Non-profits in Improving Maternal and Child Health. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36: 12017–12025.
- Modi, N.; Mary, P.; and Moy, C. 2019. Transfer restless multi-armed bandit policy for energy-efficient heterogeneous cellular network. *EURASIP Journal on Advances in Signal Processing*, 2019(1): 1–19.
- Mordatch, I.; and Abbeel, P. 2018. Emergence of grounded compositional language in multi-agent populations. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Munos, R. 2003. Error bounds for approximate policy iteration. In *Proceedings of the Twentieth International Conference on International Conference on Machine Learning*, 560–567.
- Nagaraj, D. M.; Kowshik, S. S.; Agarwal, N.; Netrapalli, P.; and Jain, P. 2023. Multi-user reinforcement learning with low rank rewards. In *International Conference on Machine Learning*, 25627–25659. PMLR.
- Nair, R.; Tambe, M.; Roth, M.; and Yokoo, M. 2004. Communications for improving policy computation in distributed POMDPs. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems, 2004. AAMAS 2004.*, 1098–1105. IEEE.
- Oliehoek, F. A.; Amato, C.; et al. 2016. *A concise introduction to decentralized POMDPs*, volume 1. Springer.
- Papadimitriou, C. H.; and Tsitsiklis, J. N. 1999. The Complexity of Optimal Queuing Network Control. *Mathematics of Operations Research*, 24(2): 293–305.
- Paulus, D.; de Vries, G.; Janssen, M.; and Van de Walle, B. 2023. Reinforcing data bias in crisis information management: The case of the Yemen humanitarian response. *International Journal of Information Management*, 72: 102663.
- Peng, B.; Rashid, T.; Schroeder de Witt, C.; Kamienny, P.-A.; Torr, P.; Böhm, W.; and Whiteson, S. 2021. Facmac: Factored multi-agent centralised policy gradients. *Advances in Neural Information Processing Systems*, 34: 12208–12221.
- Rashid, T.; Samvelyan, M.; Witt, C. S.; Farquhar, G.; Foerster, J.; and Whiteson, S. 2018. QMIX: Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. In *International Conference on Machine Learning*, 4292–4301.
- Ruiz-Hernández, D.; Pinar-Pérez, J. M.; and Delgado-Gómez, D. 2020. Multi-machine preventive maintenance scheduling with imperfect interventions: A restless bandit approach. *Computers & Operations Research*, 119: 104927.

- Schmitt-Grohé, S.; Teoh, K.; and Uribe, M. 2020. COVID-19: testing inequality in New York City. Technical report, National Bureau of Economic Research.
- Singh, A.; Jain, T.; and Sukhbaatar, S. 2019. Learning when to Communicate at Scale in Multiagent Cooperative and Competitive Tasks. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Son, K.; Kim, D.; Kang, W. J.; Hostallero, D. E.; and Yi, Y. 2019. Qtran: Learning to factorize with transformation for cooperative multi-agent reinforcement learning. In *International conference on machine learning*, 5887–5896. PMLR.
- Sukhbaatar, S.; Fergus, R.; et al. 2016. Learning multiagent communication with backpropagation. *Advances in neural information processing systems*, 29.
- Sunehag, P.; Lever, G.; Gruslys, A.; Czarniecki, W. M.; Zambaldi, V.; Jaderberg, M.; Lanctot, M.; Sonnerat, N.; Leibo, J. Z.; Tuyls, K.; et al. 2018. Value-decomposition networks for cooperative multi-agent learning based on team reward. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, 2085–2087. International Foundation for Autonomous Agents and Multiagent Systems.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement learning: An introduction*. MIT press.
- Tripathi, V.; and Modiano, E. 2019. A Whittle Index Approach to Minimizing Functions of Age of Information. In *2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 1160–1167. Allerton: Allerton.
- Verloop, I. M. 2016. Asymptotically optimal priority policies for indexable and nonindexable restless bandits. *The Annals of Applied Probability*, 26(4): 1947–1995.
- Verma, S.; Zhao, Y.; Shah, S.; Boehmer, N.; Taneja, A.; and Tambe, M. 2024. Group Fairness in Predict-Then-Optimize Settings for Restless Bandits.
- Wang, K.; Verma, S.; Mate, A.; Shah, S.; Taneja, A.; Madhiwalla, N.; Hegde, A.; and Tambe, M. 2023. Scalable decision-focused learning in restless multi-armed bandits with application to maternal and child health. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 12138–12146.
- Wang, T.; Dong, H.; Lesser, V.; and Zhang, C. 2020a. ROMA: Multi-Agent Reinforcement Learning with Emergent Roles. In *Proceedings of the 37th International Conference on Machine Learning*.
- Wang, T.; Gupta, T.; Mahajan, A.; Peng, B.; Whiteson, S.; and Zhang, C. 2020b. RODE: Learning Roles to Decompose Multi-Agent Tasks. In *International Conference on Learning Representations*.
- Wang, T.; Wang, J.; Wu, Y.; and Zhang, C. 2019a. Influence-Based Multi-Agent Exploration. In *International Conference on Learning Representations*.
- Wang, T.; Wang, J.; Zheng, C.; and Zhang, C. 2019b. Learning Nearly Decomposable Value Functions Via Communication Minimization. In *International Conference on Learning Representations*.
- Wang, T.; Zeng, L.; Dong, W.; Yang, Q.; Yu, Y.; and Zhang, C. 2021. Context-aware sparse deep coordination graphs. *arXiv preprint arXiv:2106.02886*.
- Wang, Y.; Han, B.; Wang, T.; Dong, H.; and Zhang, C. 2020c. DOP: Off-policy multi-agent decomposed policy gradients. In *International conference on learning representations*.
- Weber, R. R.; and Weiss, G. 1990. On an index policy for restless bandits. *Journal of applied probability*, 27(3): 637–648.
- Wen, M.; Kuba, J.; Lin, R.; Zhang, W.; Wen, Y.; Wang, J.; and Yang, Y. 2022. Multi-agent reinforcement learning is a sequence modeling problem. *Advances in Neural Information Processing Systems*, 35: 16509–16521.
- Whittle, P. 1988. Restless bandits: Activity allocation in a changing world. *Journal of applied probability*, 25(A): 287–298.
- Xiong, G.; and Li, J. 2023. Finite-Time Analysis of Whittle Index based Q-Learning for Restless Multi-Armed Bandits with Neural Network Function Approximation. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Yaesoubi, R.; and Cohen, T. 2011. Generalized Markov models of infectious disease spread: A novel framework for developing dynamic health policies. *European journal of operational research*, 215(3): 679–687.
- Yang, Q.; Dong, W.; Ren, Z.; Wang, J.; Wang, T.; and Zhang, C. 2022. Self-organized polynomial-time coordination graphs. In *International Conference on Machine Learning*, 24963–24979. PMLR.
- Yu, C.; Velu, A.; Vinitzky, E.; Gao, J.; Wang, Y.; Bayen, A.; and Wu, Y. 2022. The surprising effectiveness of ppo in cooperative multi-agent games. *Advances in Neural Information Processing Systems*, 35: 24611–24624.
- Zayas-Caban, G.; Jasin, S.; and Wang, G. 2019. An asymptotically optimal heuristic for general nonstationary finite-horizon restless multi-armed, multi-action bandits. *Advances in Applied Probability*, 51(3): 745–772.
- Zhao, Q.; Krishnamachari, B.; and Liu, K. 2008. On myopic sensing for multi-channel opportunistic access: structure, optimality, and performance. *IEEE Transactions on Wireless Communications*, 7(12): 5431–5440.
- Zhao, Y.; Behari, N.; Hughes, E.; Zhang, E.; Nagaraj, D.; Tuyls, K.; Taneja, A.; and Tambe, M. 2023. Towards Zero Shot Learning in Restless Multi-armed Bandits. *AAMAS*.
- Zhao, Y.; Behari, N.; Hughes, E.; Zhang, E.; Nagaraj, D.; Tuyls, K.; Taneja, A.; and Tambe, M. 2024a. Towards a Pretrained Model for Restless Bandits via Multi-arm Generalization. *IJCAI*.
- Zhao, Y.; Boehmer, N.; Taneja, A.; and Tambe, M. 2024b. Towards Foundation-model-based Multiagent System to Accelerate AI for Social Impact. *arXiv preprint arXiv:2412.07880*.
- Zhao, Y.; Pan, Q.; Choromanski, K.; Jain, D.; and Sindhwani, V. 2022. Implicit Two-Tower Policies. *arXiv preprint arXiv:2208.01191*.