

# ROMA-iQSS: An Objective Alignment Approach via State-Based Value Learning and ROUNd-ROBIN Multi-Agent Scheduling

Chi-Hui Lin, Joewie J. Koh, Alessandro Roncone, Lijun Chen

**Abstract**—Effective multi-agent collaboration is imperative for solving complex, distributed problems. In this context, two key challenges must be addressed: first, autonomously identifying optimal objectives for collective outcomes; second, aligning these objectives among agents. Traditional frameworks, often reliant on centralized learning, struggle with scalability and efficiency in large multi-agent systems. To overcome these issues, we introduce a decentralized state-based value learning algorithm that enables agents to independently discover optimal states. Furthermore, we introduce a novel mechanism for multi-agent interaction, wherein less proficient agents follow and adopt policies from more experienced ones, thereby indirectly guiding their learning process. Our theoretical analysis shows that our approach leads decentralized agents to an optimal collective policy. Empirical experiments further demonstrate that our method outperforms existing decentralized state-based and action-based value learning strategies by effectively identifying and aligning optimal objectives.

## I. INTRODUCTION

Reinforcement Learning (RL) enables autonomous agents to make informed decisions by leveraging past experiences to anticipate future rewards in a variety of contexts. However, despite its notable successes in optimizing the behavior of single- or few-agent systems, RL’s limitations (e.g. sample inefficiency, high variance, out-of-distribution behaviors) become more evident when transitioning to multi-agent settings. These challenges highlight the critical need for effective inter-agent coordination, a requirement that is particularly important for applications such as robotic warehouse management [1], [2], urban traffic light systems [3], [4], and human-robot collaboration [5], [6].

For optimal coordination, agents face two overarching challenges. First, they must identify which objectives will maximize collective utility. Second, they must achieve goal alignment with their teammates to avoid counterproductive outcomes. Traditionally, the above challenges were addressed by methods that employ a *centralized* learning approach. Specifically, these methods aggregate experiences from all agents to perform joint strategy optimization, a technique that proved effective in optimizing group outcomes [7], [8], [9], [10]. However, these approaches face three critical challenges: i) they incur exponential computational costs as the number of agents increases, ii) they necessitate the global sharing of individual agent policies, which might not be feasible in practice and limits the applicability of such




Fig. 1. Motivating example: five agents are tasked with a collaborative transport problem where a large object needs to be moved to designated locations marked by flags. In the case of failed coordination, the left subgroup (two agents), struggles to identify the optimal objectives, while the right subgroup (three agents), identifies the objectives but fails to align their efforts effectively, leading to an undesired final location for the object. In contrast, successful coordination manifests when all agents not only identify the optimal objectives but also achieve precise alignment in their efforts, culminating in the object reaching its intended location.

methods, and iii) it is impractical to maintain continuous, high-bandwidth communication with a central controller in dynamic environments, such as self-driving vehicles. Collectively, these limitations inevitably limit the scalability and robustness of centralized training paradigms.

Conversely, *decentralized* learning approaches offer a compelling alternative to centralized training, allowing each agent to estimate expected returns autonomously based on individual experiences. Crucially, these algorithms preclude direct inter-agent communication and restrict access to the policies of other agents [11], [12], [13], [14], [15]. As such, decentralized learning paradigms, such as independent Q-learning [16], are particularly well-suited for highly scalable scenarios. Nevertheless, the limited flow of information creates obstacles in agents’ understanding of their peers and their ability to interpret the environment effectively. This limitation frequently causes a lack of alignment in agents’ objectives or the selection of suboptimal goals, ultimately resulting in poor coordination.

Human coordination can thrive even without verbal communication, as silent interactions often suffice for mutual understanding. Motivated by this concept, we tackle the hurdles of limited communication within decentralized systems by examining their interaction patterns. Typically, agents operate concurrently, deriving lessons from these experiences—a method we have termed *Synchronous MultiAgent Interaction (SMA)*. While this strategy promotes efficient exploration, it also results in a continuously changing environment, complicating the agents’ capacity to align their strategies. To coun-

The authors are with the Department of Computer Science, University of Colorado, Boulder, Colorado, USA. Emails: {firstname.lastname}@colorado.edu

This work was supported by Army Research Laboratory under grant #W911NF-21-2-0126. Alessandro Roncone is with Lab0, Inc.

teract this issue, we present *ROMA*, a *ROund-Robin Multi-Agent Scheduling* protocol designed to streamline information acquisition during interaction for goal alignment. Within this framework, interaction is structured in rounds, with each round permitting only one agent to collect experience from the environment. Additionally, more knowledgeable agents utilize their better policies to guide the agent under consideration toward beneficial outcomes, while less experienced agents are afforded the opportunity for exploratory behaviors. This synergistic approach allows agents to capitalize on the expertise of their more experienced peers, culminating in aligned objectives.

Beyond aligning objectives, equipping agents with the skills to pinpoint the best goals is also crucial. To overcome the constraints of decentralized algorithms, we propose *independent QSS learning* (iQSS), an innovative approach to decentralized learning based on state values. iQSS works by assessing the value of the current global state based on potential future states, utilizing insights gained from environmental interactions. As a result, agents can more accurately anticipate the returns of possible future scenarios, guiding them more efficiently towards the optimal objectives.

In conclusion, this paper introduces a tightly integrated framework that combines independent state-based value learning, iQSS, with a specialized multi-agent interaction protocol, ROMA. These two components are not stand-alone solutions; rather, they operate synergistically to enable agents to pinpoint optimal states and synchronizing their objectives. Theoretical analysis reveals iQSS helps agents converge on effective policies for optimal states, while ROMA coordinates their efforts to a common goal. Our empirical studies, featuring multi-stage coordination tasks, demonstrate ROMA-iQSS's superiority over the current state-of-the-art, I2Q [15], and traditional independent Q learning methods. Crucially, ROMA-iQSS stands out for its dual capability to identify optimal objectives and ensure goal alignment among agents.

## II. PRELIMINARIES

### A. Coordination Problem

We aim to solve the following collaborative problem, represented by a 7-tuple Markov Decision Process,  $\{K, S, (A_k)_{k \in K}, T, r, (\pi_k)_{k \in K}, \gamma\}$ :

- $\mathcal{K} = \{1, \dots, k\}$  is the set of indexes of all agents.
- $S$  is the environment state space.
- $A_k$  is the discrete action space for agent  $k$ , while  $A_{-k}$  denotes the action space for all agents except agent  $k$ .
- $s_{(t+1)} = T(s_{(t)}, a_{(t)})$  is the environment transition function. To clarify, time step indices are indicated with parentheses, such as  $a_{(t)}$  for joint action at time  $t$ , and agent-ID indices are shown without parentheses, like  $a_k$  for agent  $k$ 's action.
- $r$  is the shared reward received by all agents  $k \in K$ .
- $\pi_k$ , belongs to a deterministic policy class  $\Pi_k$ , generates actions  $a_k \in A_k$  for agent  $k$  under environment states.
- $\gamma$  is a discount factor.

Our main goal is to develop a learning method to enable each agent to independently learn its optimal policy,  $\pi_k$ , using its own experiences, without the explicit knowledge of policies shared by other agents. More specifically, the joint of all individual optimal policies should be an optimal joint policy  $\pi^*$  that maximizes the cumulative discounted shared rewards  $r$  with a discount factor denoted as  $\gamma$ . These rewards depends on the current state,  $s_{(t)}$ , and the next state,  $s_{(t+1)}$ , arising from a deterministic transition function,  $T$ . This function uses the current state and the action derived from the policy,  $\pi(s_{(t)})$ , to generate the next state.

$$\begin{aligned} \pi^* &= \arg \max_{\pi \in \Pi} \mathbb{E} \left[ \sum_{s_{(0)} \in S} \sum_{t \geq 0} \gamma^t r(s_{(t)}, s_{(t+1)}) \right], \\ s_{(t+1)} &= T(s_{(t)}, \pi(s_{(t)})), \\ \Pi &= \Pi_1 \times \Pi_2 \times \dots \times \Pi_{|K|} \end{aligned}$$

Additionally, agents collect their own experiences when interacting with other agents in the environment. During an agent's interaction, it receives multiple pieces of information, each represented by a tuple  $(s, a_k, s', r)$ . Each tuple contains the current environment state  $s$ , an agent's own action  $a_k$ , a subsequent environment state  $s'$ , and a shared reward  $r$ .

### B. Overview of Value Learning Methods

Centralized Q learning provides a method to generate an optimal policy that maximizes a group's profits, but it suffers from scalability. On the other hand, independent Q learning enables agents to learn independently. However, it is challenging for it to learn an optimal policy due to its partial understanding of the environment.

1) *Centralized Q Learning (cenQ)*: In the context of our coordination problem, the utilization of centralized Q-learning necessitates the presence of a central controller. The central controller accesses the experience of all agents and leverages it to learn the joint strategy for all agents.

Specifically, cenQ maintains a value function, representing an expected long-term return associated with the environment state and joint action,  $a$ , denoted as  $Q(s, a)$ , where  $a = (a_1, a_2, \dots, a_{|K|})$  and  $a \in A$ , where  $A = A_1 \times A_2 \times \dots \times A_{|K|}$ . This value function is learned through a temporal difference learning process, incorporating  $\alpha$  as the learning rate,  $\gamma$  as the discount factor, and  $s'$  as the subsequent environment state, determined by the environment transition function  $T(s, a)$ .

$$\begin{aligned} Q(s, a) &\leftarrow (1 - \alpha)Q(s, a) \\ &\quad + \alpha(r + \gamma \max_{a' \in A} Q(s', a')) \end{aligned} \quad (1)$$

Temporal difference learning in eq. (1) is well-known for yielding the optimal Q-value,  $Q^*(s, a)$ , upon convergence [17].

$$Q^*(s, a) = r(s, T(s, a)) + \gamma \max_{a' \in A} Q^*(s', a')$$

With the value function  $Q^*(s, a)$ , The centralized controller can further induce the optimal policy,  $\pi^*$ :

$$\pi^*(s) = \arg \max_{a \in A} Q^*(s, a)$$

However, even with the promise of optimal convergence, the preference for cenQ diminishes when applied to extensive multi-agent systems due to inherent scalability concerns.

2) *Independent Q Learning (indQ)*: Independent Q Learning becomes a preferable option in multi-agent systems due to its superior scalability. This permits an agent to estimate its value based on individual experiences. In technical terms, each agent,  $k \in K$ , maintains the value  $Q_k$ , representing the expected return associated with its own action and environment state, through temporal difference learning:

$$Q_k(s, a_k) \leftarrow (1 - \alpha)Q_k(s, a_k) + \alpha(r + \gamma \max_{a'_k \in A_k} Q_k(s', a'_k)) \quad (2)$$

In eq. (2), the subsequent state  $s'$  is determined by the transition function,  $T(s, a)$ . However, an independent agent  $k$  is limited to accessing the information of its own action denoted as  $a_k$ , instead of the complete action set  $a$ . It can only access its action information,  $a_k$ . Therefore, the estimated value,  $Q_k(s, a_k)$ , depends on its observed transition,  $P(s'|s, a_k)$ , which is a distribution conditioned on its action and environment state.

$$Q_k^*(s, a_k) = \mathbb{E}_{s' \sim P(s'|s, a_k)} [r(s, s') + \gamma \max_{a'_k \in A_k} Q_k^*(s', a'_k)]$$

The observed transition,  $P(s'|s, a_k)$ , in an environment with deterministic transition,  $T(s, a_k, a_{-k})$  can be represented as:

$$\begin{aligned} P(s'|s, a_k) &= \sum_{a_{-k} \in A_{-k}} P(s'|s, a_k, a_{-k})P(a_{-k}|s, a_k) \\ &= \sum_{a_{-k} \in A_{-k}} \mathbb{1}_{s'=T(s, a_k, a_{-k})} P(a_{-k}|s, a_k) \\ \mathbb{1}_{s'=T(s, a_k, a_{-k})} &= \begin{cases} 1 & , \text{if } s' = T(s, a_k, a_{-k}) \\ 0 & , \text{otherwise} \end{cases} \end{aligned}$$

Therefore,  $P(s'|s, a_k)$  is dependent on the responses,  $P(a_{-k}|s, a_k)$ , from other agents. If they keep changing their response strategy, the individual observed transition can be non-stationary, making the convergence on the value function,  $Q_k(s, a_k)$ , challenging.

Furthermore, if other agents maintain stable but suboptimal strategies, the agent could learn a suboptimal value. This, in turn, results in its derived policy,  $\pi_k^*$ , also being suboptimal.

$$\pi_k^*(s) = \arg \max_{a_k \in A_k} Q_k^*(s, a_k)$$

3) *Centralized QSS Learning (cQSS)*: QSS learning distinguishes itself as a state-based value learning. Unlike action-based methods, its return estimate is not directly dependent on the actions. Instead, it is determined by only a state and a subsequent state.

When designing a centralized QSS learning, like centralized Q learning, it also necessitates including a central controller that can access the information of all agents. However, instead of estimating the state-action value, it approximates the state-state value,  $Q(s, s')$ , an expected return associated

with a state  $s \in S$ , and a neighboring state  $s' \in N(s)$ , where  $N(s)$  is a set of possible subsequent states, observed by agents during the interaction.

$$\begin{aligned} N(s) &= \{s' | P(s'|s) > 0\} \\ &= \{s' | \mathbb{1}_{s'=T(s, a)} > 0, \forall a \in A\} \end{aligned}$$

Through the temporal difference learning, the central controller learns the converged state-state value  $Q^*(s, s')$  [18].

$$Q(s, s') \leftarrow (1 - \alpha)Q(s, s') + \alpha(r + \gamma \max_{s'' \in N(s')} Q(s', s''))$$

$$Q^*(s, s') = r(s, s') + \gamma \max_{s'' \in N(s')} Q^*(s', s'')$$

Edwards et al. show the equivalence between  $Q^*(s, a)$  and  $Q^*(s, s')$ , where the subsequent state  $s'$  is determined by the transition  $T(s, a)$  [18]:

$$Q^{s^{**}}(s, a) = Q^*(s, T(s, a))$$

We can utilize the equivalence to let the central controller induce its optimal state-action value  $Q^{s^{**}}(s, a)$  and further induce the policy  $\pi_{s^{**}}$ , which is guaranteed to be optimal.

$$\pi^{s^{**}}(s) = \arg \max_{a \in A} Q^{s^{**}}(s, a)$$

### III. METHODS

#### A. Independent QSS Value Learning (iQSS)

Although centralized QSS learning can learn the optimal policy, similar to other centralized approaches, it is also not a scalable choice. Therefore, we develop an independent state-based value learning method, independent QSS learning.

While centralized QSS learning enables a central controller to learn the state-state value, independent QSS learning enables each agent to learn its own value. Each agent learns the value through the temporal difference learning, used by the centralized QSS controller. Due to the convergence property [17], [18], each agent can ultimately learn the unique optimal state-state value  $Q^*(s, s')$ . Nevertheless, since agents learn the value independently, we represent their values, respectively, with an index  $k$ .

$$Q_k(s, s) \leftarrow (1 - \alpha)Q_k(s, s) + \alpha(r + \gamma \max_{s'' \in N(s')} Q_k(s', s'')) \quad (3)$$

$$\begin{aligned} Q^*(s, s') &= Q_k^*(s, s') \\ &= r(s, s') + \gamma \max_{s'' \in N(s')} Q_k^*(s', s'') \end{aligned} \quad (4)$$

Like cQSS inducing its action-based value, iQSS agents can also utilize their states-based value to induce their action-based value function. However, unlike the centralized approach, an independent agent cannot access the information of actions, made by others. Therefore, an iQSS agent maintains a value,  $Q_k^{s^{**}}(s, a_k)$  that takes only its own action and the environment state. It induces its state-action value  $Q_k^{s^{**}}(s, a_k)$  by maximizing the state-state value over a set of

subsequent states  $N(s, a_k)$ , which is observed and recorded during its experience collection stage.

$$Q_k^{ssa^*}(s, a_k) = \max_{s' \in N(s, a_k)} Q_k^*(s, s') \quad (5)$$

$$N(s, a_k) = \{s' | P(s'|s, a_k) > 0\} \quad (6)$$

With the converged value, iQSS agents can then induce their policies:

$$\pi_k^{ss^*}(s) = \arg \max_{a_k \in A_k} Q_k^{ssa^*}(s, a_k)$$

Although the individual policies  $\pi_k^{ss^*}(s)$  were induced without considering the action information of other agents, we show that the joint policies still have the equivalence to the optimal policy  $\pi^*(s)$  under the assumption in Theorem 3.1. We later show the assumption is true in Lemma 3.2.

*Theorem 3.1:*  $\pi_k^{ss^*}$  is optimal under the set equivalence assumption:

$$\begin{aligned} & \{s' | \exists a_{-k} \in A_{-k}, \mathbb{1}_{s'=T(s, a_k, a_{-k})} > 0\} \\ & = \{s' | P(s'|s, a_k) > 0\} \end{aligned}$$

*Proof:*

$$\begin{aligned} & Q_k^{ssa^*}(s, a_k) \\ \text{by eq. (5)} & = \max_{s' \in N(s, a_k)} Q_k^*(s, s') \\ \text{by eq. (4)} & = \max_{s' \in N(s, a_k)} Q^*(s, s') \\ \text{by eq. (6)} & = \max_{s' \in \{s' | P(s'|s, a_k) > 0\}} Q^*(s, s') \\ & = \max_{s' \in \{s' | \exists a_{-k} \in A_{-k}, \mathbb{1}_{s'=T(s, a_k, a_{-k})} > 0\}} Q^*(s, s') \\ & = \max_{a_{-k} \in A_{-k}} Q^*(s, T(s, a_k, a_{-k})) \end{aligned}$$

The first three equivalence come from the definitions. The fourth is due to the assumption, mentioned in the Theorem 3.1. For the fifth equivalence, the upper-hand side maximizes state-state value over all subsequent states, which any joint actions containing  $a_k$  could, given the current state  $s$ . The lower-hand side basically does the same thing, but it brings the transition function  $T(s, a_k, a_{-k})$  into the state-state value function. Therefore, instead of maximizing over states, it maximizes over all possible joint actions while fixing the action made by agent  $k$  to be  $a_k$ .

Let  $a_k^* = \pi_k^{ss^*}(s)$ .

$$\begin{aligned} & Q_k^{ssa^*}(s, a_k^*) \\ & = \max_{a_{-k} \in A_{-k}} Q^*(s, T(s, a_k^*, a_{-k})) \\ & = \max_{a_k \in A_k} \max_{a_{-k} \in A_{-k}} Q^*(s, T(s, a_k, a_{-k})) \\ & = \max_{(a_k, a_{-k}) \in A_k \times A_{-k}} Q^*(s, T(s, a_k, a_{-k})) \\ & \geq Q^*(s, T(s, a)), \forall a \in A \end{aligned}$$

The first equivalence comes from the final equivalence of the process, we have shown. The second equivalence is from the  $\pi_k^{ss^*}(s)$  definition. For the third equivalence, the upper-hand side first maximizes the value over joint actions containing

$a_k$  for each  $a_k \in A_k$  and then maximizes it over  $a_k \in A_k$ . It considers all the Cartesian product's joint actions,  $A_k \times A_{-k}$ . The lower-hand side of the fourth equivalence maximizes over the Cartesian product of the action sets. Therefore, the fourth equivalence holds. The final equivalence is because of the equivalence  $A = A_k \times A_{-k}$ .

The inequality shows the action generated by the  $\pi_k^{ss^*}(s)$  is the optimal action when the current state is  $s$ . The property holds for all states  $s \in S$ . Therefore,  $\pi_k^{ss^*}(s)$  is optimal. ■

Since Theorem 3.1 requires the assumption, we show it is true under a condition in the Lemma 3.2. In addition, we later create multi-agent interaction schemes to let the condition hold.

*Lemma 3.2:* The set equivalence assumption, mentioned in the Theorem 3.1, holds true under the relation assumption:

$$\exists a_{-k} \in A_{-k}, \mathbb{1}_{s'=T(s, a_k, a_{-k})} > 0 \Rightarrow P(s'|s, a_k) > 0$$

*Proof:* We show the equivalence by showing they are subsets of each other. For simplicity of clarification, we represent two sets by  $X$  and  $Y$ . We first prove the set  $X$  is a subset of the set  $Y$ .

$$\begin{aligned} X & = \{s' | P(s'|s, a_k) > 0\} \\ & = \{s' | \sum_{a_{-k} \in A_{-k}} P(s'|s, a_k, a_{-k}) P(a_{-k}|s, a_k) > 0\} \\ & = \{s' | \sum_{a_{-k} \in A_{-k}} \mathbb{1}_{s'=T(s, a_k, a_{-k})} P(a_{-k}|s, a_k) > 0\} \\ & \subseteq \{s' | \sum_{a_{-k} \in A_{-k}} \mathbb{1}_{s'=T(s, a_k, a_{-k})} > 0\} \\ & = \{s' | \exists a_{-k} \in A_{-k}, \mathbb{1}_{s'=T(s, a_k, a_{-k})} > 0\} = Y \end{aligned}$$

By the law of the total probability, we have the first equivalence. Then, since the transition,  $P(s'|s, a_k, a_{-k})$ , is deterministic in our environments and represented by  $s' = T(s'|s, a_k, a_{-k})$ , we have the second equivalence. In addition, we have the subset relation between the third and the fourth line because of the inequality for any  $s' \in S$ :

$$\begin{aligned} & \sum_{a_{-k} \in A_{-k}} \mathbb{1}_{s'=T(s, a_k, a_{-k})} P(a_{-k}|s, a_k) \leq \\ & \sum_{a_{-k} \in A_{-k}} \mathbb{1}_{s'=T(s, a_k, a_{-k})} \end{aligned}$$

The inequality is true because  $0 \leq P(a_{-k}|s, a_k) \leq 1$ . Finally, the final equivalence comes from the relation:

$$\begin{aligned} & \sum_{a_{-k} \in A_{-k}} \mathbb{1}_{s'=T(s, a_k, a_{-k})} > 0 \Leftrightarrow \\ & \exists a_{-k} \in A_{-k}, \mathbb{1}_{s'=T(s, a_k, a_{-k})} > 0 \end{aligned}$$

It is true because  $\mathbb{1}$  is non-negative.

Moreover, we can prove the set  $Y$  is a subset of the set  $X$  by the assumption, mentioned in the theorem.

$$\begin{aligned} Y & = \{s' | \exists a_{-k} \in A_{-k}, \mathbb{1}_{s'=T(s, a_k, a_{-k})} > 0\} \\ & \subseteq \{s' | P(s'|s, a_k) > 0\} = X \end{aligned}$$

Since they are subsets of each, we have the set equivalence in Lemma 3.2. ■

*Lemma 3.3:* The relation assumption, mentioned in Lemma 3.2, holds true if  $P(a_{-k}|s, a_k)$  is always positive.

*Proof:* From the law of total probability, we can infer the equalities:

$$\begin{aligned} P(s'|s, a_k) &= \sum_{a_{-k} \in A_{-k}} P(s'|s, a_k, a_{-k})P(a_{-k}|s, a_k) \\ &= \sum_{a_{-k} \in A_{-k}} \mathbb{1}_{s'=T(s, a_k, a_{-k})} P(a_{-k}|s, a_k) \end{aligned}$$

Additionally, we know that  $\mathbb{1}_{s'=T(s, a_k, a_{-k})}$  is always non-negative and we assume that  $P(a_{-k}|s, a_k)$  is always positive. Therefore, the relation is true:

$$\exists a_{-k} \in A_{-k}, \mathbb{1}_{s'=T(s, a_k, a_{-k})} > 0 \Rightarrow P(s'|s, a_k) > 0$$

Based on the given Theorem and Lemmas, we deduce that the strategy  $\pi_k^{SS^*}$  is deemed optimal when the conditional probability  $P(a_{-k}|s, a_k)$  is consistently greater than zero. This assumption essentially means that for any chosen action  $a_k$  by agent  $k$ , there is always a nonzero probability of encountering any other actions  $a_{-k}$  from other agents. Put simply, regardless of the decision made by agent  $k$ , there's always a possibility to interact with the varied actions of other agents. To meet this requirement, one can implement an  $\epsilon$ -greedy strategy with  $\epsilon > 0$ , ensuring that agents occasionally choose actions at random, thereby maintaining a positive probability of diverse actions.

### B. Multiagent Interaction Schemes

Agents require experience for learning. They collect experience when interacting with other agents. In an independent learning setting, each agent collects its own experience and does not share it with others. Moreover, independent QSS learning agents require an assumption to ensure the individual learned policy  $\pi_k^{SS^*}$  to be optimal. Therefore, we design our interaction schemes to ensure it.

1) *Synchronous Multiagent Interaction (SMA):* From Theorem 3.1, Lemma 3.2 and 3.3, we know that  $\pi_k^{SS^*}$  is optimal if  $P(a_{-k}|s, a_k)$  is always positive. That means  $\pi_k^{SS^*}$  is optimal if the probability of all possible joint actions should stay positive no matter what action a certain agent executes. Therefore, to ensure  $P(a_{-k}|s, a_k) > 0$ , we need to ensure agents always take a random action with probability  $> 0$  when interacting with others. More specifically, we can enable agents to execute  $\epsilon$ -greedy policy, which takes random action with probability  $\epsilon > 0$ , and learned action with probability  $1 - \epsilon$ :

$$\pi_k^\epsilon(s) = \begin{cases} \pi_k(s) & , \text{with probability } 1 - \epsilon \\ \text{a random action} & , \text{with probability } \epsilon \end{cases}$$

In synchronous multiagent interaction (SMA), each agent executes  $\epsilon$ -greedy policy with  $\epsilon > 0$  and collects experiences to its replay buffers. Then, from its collected experience, agent  $k \in K$  would observe  $P(a_{-k}|s, a_k) > 0$ . Therefore, with Theorem 3.1, we know each agent  $k \in K$  would learn an individual policy,  $\pi_k^{SS^*}$ , which is optimal.

### Algorithm 1 Synchronuous Multiagent Interaction

- 1: **for**  $t = 1, 2, \dots, t_{max}$  **do**
- 2:   All agents execute their  $\epsilon$ -greedy policy.
- 3:   All agents record the interaction to the replay buffers.
- 4: **end for**




Fig. 2. Problem Definition: This is a 2-agent game featuring four potential destination states, each represented by two numbers denoting Agent X's and Agent Y's policies. **In this game, both agents can identify optimal states, indicated in red text, but they make their decisions independently without knowledge of each other's choices.** *SMA-Scenario(Left):* Both agents observe all potential states, including two optimal states. Consequently, they might establish divergent objectives, which prevents them from reaching either of the two optimal states. *ROMA-Scenario(Right):* Agent X observes all potential states and selects its policy accordingly. Agent Y, on the other hand, observes only the states that Agent X's selection can lead to. Consequently, Agent X's selection influences Agent Y's choice, enabling it to align with Agent X's objective and ultimately reach the optimal state.

2) *Round-Robin Multiagent Interaction (ROMA):* When all agents follow the SMA interaction process, each agent can learn an individual optimal policy through independent QSS learning. However, they might not form an optimal joint policy if there exists more than one optimal joint policy in the environment. More specifically, if agents learn their individual optimal policy aiming at reaching a divergent optimal joint policy, the joint of their policies usually ends up being policies, that are not the optimal policy, expected by them. Take a robot coordination task as an example; all robots acknowledge that the optimal solution is that they all turn right or left. Therefore, each of them will learn that the optimal individual action could be either turning left or right. A potential joint of its optimal individual policies would be that some turn left, and some turn right. That is apparently not the result generated by an optimal joint policy.

We design an interaction scheme, round-robin multiagent interaction (ROMA), for agents to collect limited, but sufficient information that facilitates them in aligning their objectives with others so as to cooperate in learning the joint optimal. In ROMA, in each iteration, only one agent can collect experience. We refer to that single agent as the collector. We let a certain agent,  $c$ , be the only collector for continuous  $t_u$  iterations. After agent  $c$ 's collection period, we let its next agent, which has an index  $(c + 1) \bmod |K|$ , be the only collector at the next continuous  $t_u$  iterations. We enable agents to repeat the rotation process.

During the iterations, where agent  $c$  is the only collector, agents with an index, smaller than  $c$  are seniors, and agents with an index, greater than  $c$ , are juniors. Moreover, seniors can only execute their learned policy to interact with others. All other agents, including all juniors and the collector,

execute their  $\epsilon$ -greedy policies, which allows them to take random actions. Consequently, the collector, agent  $c$  observes experience, caused and limited by the joint of the learned policies of seniors. In other words, the collector cannot observe the transition conditioned on the joint actions, which its seniors do not take at all.

---

**Algorithm 2** Round-Robin Multiagent Interaction

---

- 1: **for**  $t = 1, 2, \dots, t_{max}$  **do**
  - 2:      $c = \lceil \frac{t}{t_u} \rceil \bmod |K|$
  - 3:      $collector = \text{agent } c$
  - 4:      $seniors Z_c = \text{agents with an index } < c$
  - 5:      $juniors J_c = \text{agents with an index } > c$
  - 6:      $\forall z \in Z_c$  execute their learned policies,  $\forall j \in J_c$  and agent  $c$  execute their  $\epsilon$ -greedy policies
  - 7:      $collector$  records the interaction to replay buffer
  - 8: **end for**
- 

The observation for a collector is limited to the transitions caused by the joint policy of its seniors. Therefore, if seniors have coordinated in learning a joint optimal, the limited observation is still sufficient to enable the collector to observe the optimal transitions so as to learn an optimal policy. Moreover, the limited observation, caused by seniors, also facilitates the collector to align its objectives with seniors. Consequently, following ROMA, the collector can coordinate with its sensors to learn a joint optimal policy even if multiple optimal joint policies exist.

As a result, ROMA enables agents to reach optimal coordination. We place the proof of the theorem in the appendix, which can be found on ArXiv.

*Theorem 3.4:* If all agents follow the ROMA interaction process to collect experiences, the joint of independent QSS learning agents' individual converged strategies,  $\pi_k^{s*s*}$ , is optimal.

#### IV. A PRACTICAL IQSS LEARNING PROCESS

During the interaction, each agent collects its experience into its replay buffer. From the replay buffer, each agent  $k \in K$  can observe multiple experienced tuples, each of which is represented as  $(s, a_k, s', r)$  that are the subsequent state  $s'$  and the received reward  $r$  after taking the action  $a_k$  at the state  $s$ . With the data, each agent can learn its policy by modeling the transition  $P(s'|s, a_k)$  and the neighboring states  $N(s, a_k)$ , identifying the optimal states, optimizing state-based value network  $Q(s, s')$ , and finally inducing the state-action value function  $Q^{s*s*a}(s, a_k)$ , and the policy  $\pi^{s*s}(s)$ .

Each agent  $k \in K$  first learns the observed transition,  $P(s'|s, a_k)$ , by modeling it using a neural network, represented by parameters  $\theta$  and maximizing the likelihood of the observed data point.

$$P(s'|s, a_k) \approx P(s'|s, a_k, \theta)$$

An agent can use the estimated likelihood to infer the neighboring states. If the likelihood  $P(s'|s, a_k, \theta)$  is greater than  $\delta$ , that means that an agent  $k \in K$  observes a subsequent

state  $s'$  with a probability higher than  $\delta$  when it takes action  $a_k$  at the state  $s$ . Therefore, it can consider the subsequent state  $s'$  as a possible neighboring state,  $s' \in N^\delta(s, a_k, \theta)$ . Additionally, when  $\delta$  is zero, an agent considers all possible states it has observed from the collected data. It is fine if other agents do not change their policy at all. However, if others change their policies, an agent might consider states, which were possible but not going to be reached. Therefore, each agent should set  $\delta$  to be close to, but greater than zero. As a result, it would only consider states that are still possible to reach with a probability higher than  $\delta$ , greater than zero, and thus, would not consider those that will not be reached.

$$N^\delta(s, a_k, \theta) \approx \{s' | P(s'|s, a_k, \theta) > \delta\}$$

Each agent also needs to estimate the state-based value  $Q(s, s')$  through the temporal difference learning using eq. (3). When applying it, agents must maximize the state-based value over all possible neighboring states of the state. However, the space of neighboring states is usually large. Therefore, an efficient way to identify optimal neighboring states is desirable.

$$\begin{aligned} & \max_{s' \in N(s)} Q_k(s, s') \\ &= \max_{a_k \in A_k} \max_{s' \in N(s, a_k)} Q_k(s, s') \\ &= \max_{a_k \in A_k} Q_k(s', \hat{z}_{s, a_k}) \end{aligned} \quad (7)$$

We can rewrite the maximization using individual neighboring states  $N(s, a_k)$  as the second line of the eq. (7), which maximizes over an individual action space and also the individual neighboring state. The maximization over individual neighboring states  $N(s, a_k)$  is computing expensive. Therefore, we develop a method to compute the optimal individual neighboring state  $\hat{z}_{s, a_k}$ . Consequently, we can compute the maximization by only maximizing over an individual action space.

$$\hat{z}_{s, a_k} = \arg \max_{s' \in N_k(s, a_k, \theta)} Q_k(s, s')$$

The optimal individual neighboring state  $\hat{z}_{s, a_k}$  is the state, which provides the max state-based value among all individual neighboring states  $N(s, a_k)$ . We enable agents to learn  $\hat{z}_{s, a_k}$  by an iterative update method. We first assign an agent a random state as  $\hat{z}_{s, a_k}$ . However, if  $P(\hat{z}_{s, a_k} | s, a_k, \theta)$  is smaller than  $\delta$ , that means  $\hat{z}_{s, a_k}$  does not belong to  $N(s, a_k)$ . Therefore, we must update it by a neighboring state  $s'$ , which belongs to  $N(s, a_k)$ , where  $P(s' | s, a_k, \theta) > \delta$ . Additionally, if it belongs to  $N(s, a_k)$ , but offers a worse performance than the other individual neighboring state. Using the better state, we must update  $\hat{z}_{s, a_k}$ . These two basic principles enable agents to improve their understanding of  $\hat{z}_{s, a_k}$  and ultimately learn the optimal neighboring state:

$$\hat{z}_{s, a_k} \leftarrow \begin{cases} s', & \text{if } P(\hat{z}_{s, a_k} | s, a_k, \theta) \leq \delta \\ & \text{and } P(s' | s, a_k, \theta) > \delta \\ s', & \text{if } Q_k(s, \hat{z}_{s, a_k}) < Q_k(s, s') \\ & \text{and } P(s' | s, a_k, \theta) > \delta \\ \hat{z}_{s, a_k}, & \text{else} \end{cases}$$

With the optimal neighboring state, we can develop an efficient temporal difference learning method to update the state-based value  $Q_k(s, s)$  by replacing the maximization with eq. (7):

$$Q_k(s, s') \leftarrow (1 - \alpha)Q_k(s, s') + \alpha(r + \gamma \max_{a_k \in A_k} Q_k(s', \hat{z}_{s', a_k})) \quad (8)$$

The temporal difference learning in eq. (8) enables each agent to update the value  $Q^k(s, s')$ . In addition, during each iteration, they can use the states-based value  $Q(s, s')$  to update the state-action value  $Q_k^{ssa}(s, a_k)$  and also further compute the policy  $\pi_k^{ss}(s)$ :

$$Q_k^{ssa}(s, a_k) \leftarrow \max_{s' \in N(s, a_k)} Q_k(s, s')$$

$$\pi_k^{ss}(s) \leftarrow \arg \max_{a_k \in A_k} Q_k^{ssa}(s, a_k)$$

As a result, given enough time each agent learns optimal values and the optimal policy. If an overview of practical iQSS learning is desired, we provide the pseudocode in the appendix, which can be found on ArXiv.

## V. NUMERICAL EXPERIMENTS

We evaluate our algorithms via a coordination challenge that involves independent agents. Every agent begins with knowledge of various strategies but does not know which strategy offers the most group benefit and what strategies other agents are considering. The goal is for agents to tackle two challenges: to independently identify the best strategies and to align their efforts with other agents.

The issue is divided into three phases, where agents must address two challenges in each to increase collective earnings. Not pinpointing the best strategies leads to low rewards, ranging from 0 to 10. Conversely, agents targeting divergent optimal states face hefty fines of  $-200$ . Achieving success hinges on all agents aligning their goals and collaborating towards common objectives, which can yield significant gains of 100.

Agents independently choose from three actions at each stage, forming a joint strategy. Notably, only two specific joint strategies yield successful coordination, directing agents toward one of the two optimal states. Furthermore, successful coordination hinges on all agents converging to pursue the same optimal objective; otherwise, effective coordination is unlikely.

We simulate coordination among three, five, and seven agents. For three agents, the task is to choose the best state out of 65 states. For five and seven agents, the challenge increases with 100 states. Our investigation compares our novel state-based Q learning technique (iQSS) against the established state-based Q learning method (I2Q) and the conventional action-based Q learning (independent Q Learning). iQSS and I2Q both rely on state values for making decisions but differ in their learning approaches and strategies for selecting the optimal state.

We assess two interaction strategies: SMA (standard) and ROMA (innovative), applying them across all examined

learning approaches. This results in six distinct evaluative comparisons: ROMA-iQSS, SMA-iQSS, ROMA-I2Q, SMA-I2Q, and the independent Q learning techniques including ROMA-indQ and SMA-indQ. We carry out fifteen simulations for each method across scenarios involving 3, 5, and 7 agents to analyze the average group reward and performance indicators. Furthermore, we choose an  $\epsilon$ -greedy policy setting of  $\epsilon = 0.8$  to encourage comprehensive exploration. This high  $\epsilon$  value is crucial for simulating the objective alignment challenge, ensuring that the majority of agents experience both optimal states. Additionally, adopting a higher  $\epsilon$  value helps in minimizing the likelihood of agents becoming entrapped in local optima. Moreover, we show more experiment setup details in the appendix, which can be found on ArXiv.

The analysis reveals that ROMA significantly diminishes performance variance compared to SMA algorithms. Specifically, in scenarios with three and five agents, the performance of SMA-iQSS demonstrates considerable fluctuations, whereas ROMA-iQSS stabilizes quickly, achieving consistent and superior performance after a few iterations. Furthermore, both ROMA-indQ and ROMA-I2Q exhibit marginally better and notably steadier performance than their SMA equivalents. This reduction in variance can be attributed to the more stationary environment created by the ROMA structure. In SMA settings, all agents are concurrently exploring and making decisions to achieve their goals, leading to challenges in objective alignment when goals differ, thereby producing a non-stationary environment. Conversely, ROMA facilitates a more stable environment by enabling agents to naturally follow the lead of more experienced or "senior" agents in decision-making processes, thereby streamlining objective alignment without the need for explicit discussions.

The plots demonstrate that iQSS, particularly ROMA-iQSS, outperforms I2Q and independent Q learning in efficiently finding optimal states. ROMA-iQSS consistently achieves superior results, indicating its effectiveness in identifying the best state. Although SMA-iQSS does not always deliver top average performance, it reaches significantly higher peaks in some iterations, showing it can find the optimal state despite challenges in aligning objectives. In contrast, I2Q and independent Q learning typically do not attain the significant rewards, emphasizing the advantage of iQSS in state optimization. This advantage stems from I2Q's differing method of state estimation and the limitations of action-based techniques like independent Q learning in achieving similar outcomes.

## VI. CONCLUSION AND DISCUSSION

In this work, we present a novel decentralized framework decentralized framework aimed at overcoming the challenge of forming an optimal collective strategy without direct communication. To address this challenge, we develop two critical components. Firstly, we develop an independent states-based learning method, iQSS, designed to facilitate each agent to efficiently identify optimal states and individual policies. Secondly, we introduce an interaction scheme, ROMA, which empowers agents to collect limited




Fig. 3. Teams of 3, 5, and 7 agents navigate three-stage coordination, targeting top outcomes in an environment with multiple optimal strategies.

yet sufficient data. This data allows each agent to implicitly guide less experienced peers in their learning journey while simultaneously benefiting from the guidance of more senior agents during their own learning process. As a result, ROMA plays a pivotal role in aligning the objectives of all agents, ultimately leading to successful coordination.

We substantiate our claims with rigorous proof, affirming that iQSS is capable of learning individual optimal policies, while also demonstrating that ROMA fosters objective alignment among agents. Furthermore, we provide empirical evidence through multi-agent coordination simulations, showcasing our method’s superiority over the current state-of-the-art state-based value learning approach, I2Q, and the classical action-based, independent Q learning method.

For future work, we aim to extend our multi-agent coordination methods to real-world applications, such as coordinating robots in warehouses or ensuring safe autonomous driving for automobiles. Additionally, we plan to incorporate human-like features into our agents. While human strategies are often inaccessible to robots, our framework equips agents to coordinate effectively with others even without insights into their strategies. Consequently, we are optimistic that our

approach holds significant potential to enhance human-robot coordination.

## REFERENCES

- [1] G. Papoudakis, F. Christianos, L. Schäfer, and S. V. Albrecht, “Benchmarking Multi-Agent Deep Reinforcement Learning Algorithms in Cooperative Tasks,” in *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks (NeurIPS)*, 2021. [Online]. Available: <http://arxiv.org/abs/2006.07869>
- [2] J. J. Koh, G. Ding, C. Heckman, L. Chen, and A. Roncone, “Cooperative Control of Mobile Robots with Stackelberg Learning,” in *Proceedings of the 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 7985–7992.
- [3] X. Wang, L. Ke, Z. Qiao, and X. Chai, “Large-scale Traffic Signal Control Using a Novel Multiagent Reinforcement Learning,” *IEEE Transactions on Cybernetics*, vol. 51, no. 1, pp. 174–187, 2021.
- [4] H. Wei, N. Xu, H. Zhang, G. Zheng, X. Zang, C. Chen, W. Zhang, Y. Zhu, K. Xu, and Z. Li, “CoLight: Learning Network-Level Cooperation for Traffic Signal Control,” in *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. ACM, Nov 2019.
- [5] J. Brawer, D. Ghose, K. Candon, M. Qin, A. Roncone, M. Vázquez, and B. Scassellati, “Interactive Policy Shaping for Human-Robot Collaboration with Transparent Matrix Overlays,” in *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, 2023, pp. 525–533.
- [6] Y.-S. Tung, M. B. Luebbbers, A. Roncone, and B. Hayes, “Workspace Optimization Techniques to Improve Prediction of Human Motion During Human-Robot Collaboration,” in *In Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction (HRI’24)*, ACM/IEEE. New York, NY, USA: ACM, Mar 2024. [Online]. Available: <https://doi.org/10.1145/3610977.3635003>
- [7] K. Son, D. Kim, W. J. Kang, D. E. Hostallero, and Y. Yi, “QTRAN: Learning to factorize with transformation for cooperative multi-agent reinforcement learning,” in *International conference on machine learning*. PMLR, 2019, pp. 5887–5896.
- [8] J. K. Gupta, M. Egorov, and M. Kochenderfer, “Cooperative Multi-Agent Control using Deep Reinforcement Learning,” in *Autonomous Agents and Multiagent Systems: AAMAS 2017 Workshops, Best Papers, Revised Selected Papers 16*. Springer, 2017, pp. 66–83.
- [9] R. Lowe, Y. Wu, A. Tamar, J. Harb, P. Abbeel, and I. Mordatch, “Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments,” in *Advances in Neural Information Processing Systems*, 2017.
- [10] S. Iqbal and F. Sha, “Actor-Attention-Critic for Multi-Agent Reinforcement Learning,” in *Proceedings of the International Conference on Machine Learning*, 2019.
- [11] M. Tan, “Multi-agent Reinforcement Learning: Independent vs. cooperative Agents,” in *Proceedings of the tenth international conference on machine learning*, 1993, pp. 330–337.
- [12] C. S. de Witt, T. Gupta, D. Makoviychuk, V. Makoviychuk, P. H. S. Torr, M. Sun, and S. Whiteson, “Is Independent Learning All You Need in the StarCraft Multi-Agent Challenge?” 2020.
- [13] G. Palmer, K. Tuyls, D. Bloembergen, and R. Savani, “Lenient Multi-Agent Deep Reinforcement Learning,” in *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2018.
- [14] L. Matignon, G. J. Laurent, and N. Le Fort-Piat, “Hysteretic Q-learning: An algorithm for Decentralized Reinforcement Learning in Cooperative Multi-Agent Teams,” in *Proceedings of the 2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2007.
- [15] J. Jiang and Z. Lu, “I2Q: A Fully Decentralized Q-Learning Algorithm,” *Thirty-Sixth Annual Conference on Neural Information Processing Systems*, 2022.
- [16] M. Tan, “Multi-Agent Reinforcement Learning: Independent vs. Cooperative Learning,” *Readings in Agents*, pp. 487–494, 1997.
- [17] C. J. Watkins and P. Dayan, “Q-learning,” *Machine Learning*, vol. 8, pp. 279–292, 1992.
- [18] A. Edwards, H. Sahni, R. Liu, J. Hung, A. Jain, R. Wang, A. Ecoffet, T. Miconi, C. Isbell, and J. Yosinski, “Estimating Q(s,s’) with Deep Deterministic Dynamics Gradients,” in *Proceedings of the 37th International Conference on Machine Learning*, 2020, pp. 2825–2835.