

Enhancing Multi-Agent Coordination through Common Operating Picture Integration

Peihong Yu Boram Lee Aswin Raghavan Supun Samarasekara
SRI International
first.last@sri.com

Pratap Tokekar
University of Maryland
tokekar@umd.edu

James Zachary Hare
DEVCOM Army Research Laboratory
james.z.hare.civ@army.mil

Abstract: In multi-agent systems, agents possess only local observations of the environment. Communication between teammates becomes crucial for enhancing coordination. Past research has primarily focused on encoding local information into embedding messages which are unintelligible to humans. We find that using these messages in agent’s policy learning leads to brittle policies when tested on out-of-distribution initial states. We present an approach to multi-agent coordination, where each agent is equipped with the capability to integrate its (history of) observations, actions and messages received into a Common Operating Picture (COP) and disseminate the COP. This process takes into account the dynamic nature of the environment and the shared mission. We conducted experiments in the StarCraft2 environment to validate our approach. Our results demonstrate the efficacy of COP integration, and show that COP-based training leads to robust policies compared to state-of-the-art Multi-Agent Reinforcement Learning (MARL) methods when faced with out-of-distribution initial states.

Keywords: Multi-Agent RL, Learn to Communicate, Common Operating Picture

1 Introduction

We consider multi-agent tasks where agents are required to collaborate to accomplish a shared mission. Multi-Agent Reinforcement Learning (MARL) algorithms can learn decentralized policies, such that after training, each agent uses only its local observations. In such a setting, sharing information with each other can lead to better decision-making. In this paper, we present a technique that allows agents to *learn* to communicate in a unified representation grounded in observable quantities.

Specifically, we equip each agent to integrate its (history of) own observations, actions and the communication messages into a Common Operating Picture (COP) [1]. In this short paper, we consider the state of the underlying DecPOMDP as the COP. The COP contains the position and attributes of all the agents. This COP is interpretable by human operators, making it valuable when considering human involvement in the decision-making process. We hypothesize that it improves the policy learning for each agent and essentially reduces the problem to single-agent RL. With better situational-awareness including non-local observations, we show that the COP leads to out-of-distribution generalization.




Figure 1: The purple and blue agents observe each other, the orange agent is beyond visual range. This leads to multiple potential COPs.

There are several challenges in integrating observations and messages to form a COP in a decentralized fashion. When each agent operates in an egocentric frame of reference, it is difficult to integrate and align local observations (Figure 1). Agents are often constrained by a limited range of communication, introducing uncertainty into the COP-building process. This limitation can hinder the transmission of critical information. It is impractical or unsafe to communicate raw observations. When agents communicate encoded COPs, the decoding process may introduce errors, potentially leading to hallucinations or incorrect interpretation of the communication.

A quick experiment reveals that inferring such a COP will significantly benefit the QMIX [2] family of MARL algorithms. Within the Starcraft2 multi-agent challenge (SMAC) [3], we tested QMIX on initial states (Test1-3) different from trained initial states. The experiment showed brittleness of QMIX [2] (without communication) and NDQ [4] (with inter-agent communication). When we allow QMIX to access the state (from the simulator), the learned policy is less brittle as expected. This observation suggests that transmitting messages that help infer the global state can improve generalization in OOD scenarios. We present a MARL method that grounds the learned communication in the state.




Figure 2: Motivating COPs in MARL.

2 Method

We consider the fully cooperative MARL problem with communication, which can be modeled as Decentralized Partially Observable Markov Decision Process (Dec-POMDP) [5] and formulated as a tuple $\langle \mathcal{N}, \mathcal{S}, \mathcal{A}, P, \Omega, O, R, \gamma, \mathcal{C} \rangle$. The sets $\mathcal{N} = \{1, \dots, n\}$ denotes the agents, \mathcal{S} are the states, \mathcal{A} are the actions, Ω are all observations, and \mathcal{C} are all possible communication messages. Each agent $i \in \mathcal{N}$ acquires an observation $o_i \in \Omega, o_i = O(s, i), s \in \mathcal{S}$. A joint action $\mathbf{a} = \langle a_1, \dots, a_n \rangle$ leads to the next state $s' \sim P(s'|s, \mathbf{a})$ and a shared global reward $r = R(s, \mathbf{a})$.

Each agent selects action based on observation-action history $\tau_i \in \mathcal{T} \equiv (\Omega \times \mathcal{A})^*$ using a policy $\pi(a_i|\tau_i, c_i^{in})$. The policy is shared across agents during training. Incoming communication for agent i , $c_i^{in} = [c_j^{out} \in \mathcal{C} \text{ if } d(i, j) < \rho]$ is determined by the communication range ρ and distance between agents $d(i, j)$. The overall objective is to find a joint policy $\pi(\tau, \mathbf{a})$ to maximize the global value function $Q_{tot}^\pi(\tau, \mathbf{a}) = \mathbb{E}_{s, \mathbf{a}} [\sum_{t=0}^{\infty} \gamma^t R(s, \mathbf{a}) | s_0 = s, \mathbf{a}_0 = \mathbf{a}, \pi]$, where τ is the joint observation-action history and $\gamma \in [0, 1)$ is the discount factor. We follow the *Centralized Training and Decentralized Execution* (CTDE) paradigm and adopt QMIX’s [2] architecture to build our algorithm, as illustrated in Figure 4.

We make a key assumption over prior work: agents have access to the initial state s_0 . Our approach “tracks” the COP as a transformation of the initial state informed by observation-action history and incoming communication. We communicate two embeddings — local observation-action embeddings and embeddings describing state evolution. The COP Formation module involves autoencoding these two key components: (1) In the LOP (local operating picture) thread, we encode the local observation o into an embedding denoted as z_{lop} and then decode o' . (2) In the COP thread, we encode all communicated messages using a Self-Attention network to deal with varying numbers of communication inputs, and obtain an embedding z_{cop} .

The state is tracked by GRU in hidden vector h_{cop} using the initial state. Subsequently, h_{cop} is decoded to produce the state s' . Note that \hat{s}_i^0, \hat{s}_i^t are egocentric states centered on agent i . We use \hat{s}_i^0 to initialize the GRU. The final COP is produced by




Figure 3: Masking in the formation of egocentric COP. v is the visual range of the agent and m is the map size.




Figure 4: The pipeline of our method. (a) The overall architecture. (b) Information integration and policy learning. (c) COP formation module.

combining decoded LOP o' and s' as $\hat{s}_i^t = \hat{s}_i^t \times \mathbf{1}_{v_i} + o_i^t \times (1 - \mathbf{1}_{v_i})$, where $\mathbf{1}_{v_i}$ is a mask based on visual range of agent i . The masking is illustrated in Figure 3.

Our CTDE training objective comprises three parts $\mathcal{L} = \mathcal{L}_{RL} + \mathcal{L}_{lop} + \mathcal{L}_{cop}$, where (1) \mathcal{L}_{RL} is the Temporal Difference (TD)-based loss of QMIX, (2) $\mathcal{L}_{lop} = L_{mse}(o', o)$ is the local observation reconstruction loss averaged over agents and time, and (3) $\mathcal{L}_{cop} = L_{mse}(\hat{s}', \hat{s})$ is the egocentric state reconstruction loss averaged over agents and time. It is worth noting that COP formation can also be trained under a fixed policy. By producing an egocentric COP, the agent does not depend on knowledge of its own location in global frame. We produce an egocentric COP for each agent, and by comparing COPs across agents one could quantify uncertainty across agents.

3 Experiments

Methods: We compare OOD generalization of our method against: (1) MASIA [6] predicts the state from other agents' observations within a QMIX method, (2) NDQ [4] and TarMAC [7] are prior work on MARL where the learned communication is not grounded on state prediction, (3) QMIX [2] and QMIX-based baselines QMIX_wState and QMIX-Att where each agent has access to the state and other agents' observation, respectively. QMIX_wState uses the state vector in addition to the agents' observation, a strong baseline for our method that assumes known initial state s_0 . QMIX-Att uses an attention network to integrate others' observations, baseline for the self-attention architecture. Baseline architectures are shown in the appendix.

Scenarios: We evaluate these methods on three maps from the Starcraft Multi-Agent Challenge [3] namely *1o_10b_vs_1r*, *1o_2r_vs_4r*, and *3s_vs_5z*. The first two maps involve an aerial agent to track and communicate the location of an enemy agent to other friendly agents. The third map is a micromanagement task. For each map, we created OOD initial states that are different than the training initial states. The OOD states are shown in the appendix. The methods are trained on the training map and evaluated in terms of win rate using episodes with the OOD initial states.

Improved Generalization: As shown in Figure 5, our method achieves significantly higher win rates on all three maps and in 8 out of 9 OOD states tested, compared to QMIX baselines and the prior MARL work. Compared to QMIX_wState, the prior work MASIA, NDQ and TarMAC shows brittleness. Our method performs even better than QMIX_wState, possibly due to training embeddings that produce both winning policies and accurate state prediction.




Figure 5: Comparison of methods on OOD initial states




Figure 6: Training curves of various methods

Sample Complexity: We observe (Figure 6) that our method has significantly lower sample complexity and faster convergence on the maps that require communication beyond visual range (*1o_10b_vs_1r* and *1o_2r_vs_4r*). Obviously, the performance of QMIX plateaus due to lack of communication, and MASIA and NDQ converge at a slower rate. These plots indicate that state prediction as in our COP-based training is an effective bias for quickly converging to robust policies.

Compression: To understand the success of our method further, we reduce the embedding size from 32 to 16, 8, then 4 (Figure 7). At lower message length, the win rate on OOD maps is surprisingly only a bit lower. Even at message length of 4, our method produces a significantly higher win rate than all the comparison methods. In the Appendix, we further ablate the impact of communication by dropping all messages in test episodes.

Evaluation of COPs: Figure 8 shows the average MSE of the produced COPs on training and OOD maps. We find that the COPs are highly accurate on training and the first two OOD maps (an MSE of 0.05 is roughly only 6 pixels in a field of view of 9 pixels). We see clearly that the accuracy of the MSE (in our method) is directly correlated and predictive of win rate on OOD states. Beyond OOD generalization, our method produces human-interpretable COPs. Figure 9 shows an example COP where the location of the enemy is correctly predicted despite OOD initial state.




Figure 7: Varying message size (*1o_10b_vs_1r*).




Figure 8: MSE vs win rate (*1o_2r_vs_4r*).




Figure 9: An example of reconstructed COP (*1o_10b_vs_1r*).

Acknowledgments

This material is based upon work supported by the Army Research Laboratory (ARL) under the Army AI Innovation Institute (A2I2) program Contract No. W911NF-22-2-0137. This work is supported in part by ONR under grant number N00014-18-1-2829.

References

- [1] K. Steen-Tveit and B. Erik Munkvold. From common operational picture to common situational understanding: An analysis based on practitioner perspectives. *Safety Science*, 142:105381, 2021. ISSN 0925-7535. doi:<https://doi.org/10.1016/j.ssci.2021.105381>. URL <https://www.sciencedirect.com/science/article/pii/S0925753521002253>.
- [2] T. Rashid, M. Samvelyan, C. S. de Witt, G. Farquhar, J. Foerster, and S. Whiteson. Qmix: Monotonic value function factorisation for deep multi-agent reinforcement learning. arxiv e-prints, page. *arXiv preprint arXiv:1803.11485*, 2018.
- [3] O. Vinyals, T. Ewalds, S. Bartunov, P. Georgiev, A. S. Vezhnevets, M. Yeo, A. Makhzani, H. Küttler, J. Agapiou, J. Schrittwieser, et al. Starcraft ii: A new challenge for reinforcement learning. *arXiv preprint arXiv:1708.04782*, 2017.
- [4] T. Wang, J. Wang, C. Zheng, and C. Zhang. Learning nearly decomposable value functions via communication minimization. *arXiv preprint arXiv:1910.05366*, 2019.
- [5] F. A. Oliehoek, C. Amato, et al. *A concise introduction to decentralized POMDPs*, volume 1. Springer, 2016.
- [6] C. Guan, F. Chen, L. Yuan, C. Wang, H. Yin, Z. Zhang, and Y. Yu. Efficient multi-agent communication via self-supervised information aggregation. *Advances in Neural Information Processing Systems*, 35:1020–1033, 2022.
- [7] A. Das, T. Gervet, J. Romoff, D. Batra, D. Parikh, M. Rabbat, and J. Pineau. Tarmac: Targeted multi-agent communication. In *International Conference on Machine Learning*, pages 1538–1546. PMLR, 2019.
- [8] C. Zhu, M. Dastani, and S. Wang. A survey of multi-agent reinforcement learning with communication. *arXiv preprint arXiv:2203.08975*, 2022.
- [9] T. Lin, J. Huh, C. Stauffer, S. N. Lim, and P. Isola. Learning to ground multi-agent communication with autoencoders. *Advances in Neural Information Processing Systems*, 34:15230–15242, 2021.
- [10] J. Shao, H. Zhang, Y. Qu, C. Liu, S. He, Y. Jiang, and X. Ji. Complementary attention for multi-agent reinforcement learning. 2023.

A Additional Experiments

A.1 Related Works

Communication plays a vital role in multi-agent reinforcement Learning (MARL) systems [8], enabling agents to coordinate, share information, and collectively solve complex tasks. However, there remains an open question regarding what should be communicated and how to effectively utilize these communication messages. In various approaches, such as TarMAC [7], the sending agents generate both a signature and a message derived from their local observations. Meanwhile, the receiving agents employ attention mechanisms to combine and integrate these messages based on their relevance. Other methods, like NDQ [4], introduce information-theoretic regularizers to minimize overall communication while maximizing the information conveyed by the messages for enhanced coordination. Some recent works explore more interpretable communications. For instance, [9] aims to establish a common understanding in communication symbols by autoencoding raw observations as the communication messages. Additionally, MASIA [6] adopts an approach where raw observations are communicated and aggregated into a latent representation grounding the true state, aiming to reconstruct the global state from this embedding.

A.2 Motivating Example in Details

Since the COP leads to better situational-awareness, we hypothesize that it also leads to improved out-of-distribution generalizations. We illustrate this with the following example. We use the *1o_10b_vs_1r* map from the Starcraft game [3]. This map comprises a team consisting of 1 flying unit (Overseer) and 10 ground units (Banelings), with the collective objective of eliminating a single enemy unit (Roach). The flying and enemy units are initialized simultaneously at one of the four map corners. The ten ground units are distributed randomly across the map.

To introduce diversity and challenge, we have modified the initial distribution of the flying and enemy units and incorporated three out-of-distribution scenarios for this map, as shown in Figure 11a. In Test 1, the setup closely resembles the training scenario. After selecting one corner to generate the enemy unit, we randomly position the flying unit within the nearby circle corresponding to that corner. Test 2 and Test 3 present greater challenges with notably different initial state distributions. In Test 2, flying and enemy units are generated together across the entire map, introducing a more complex spatial arrangement. In Test 3, the flying unit is generated inside the middle circle, considerably distant from the enemy unit, creating a scenario demanding different strategic considerations.

We want to answer the following question: *Will policies exhibit improved generalization on out-of-distribution (OOD) laydowns when agents possess a higher degree of situational awareness?* We employed two strong baseline algorithms to explore this question: QMIX [2] (without communication) and NDQ [4] (with inter-agent communication). Additionally, we implemented QMIX-wstate, which does not utilize communication but assumes that each agent has access to both its own partial observation and the global state. As shown in Figure 10, our findings indicate that QMIX-wstate outperforms the other algorithms across all OOD laydown scenarios. This observation strongly suggests that using the global state, i.e., COP, is advantageous for enhancing generalization in OOD scenarios.



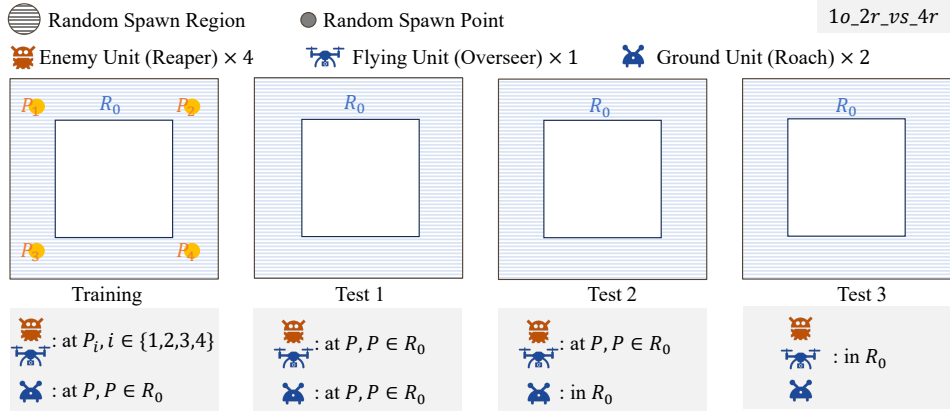


Figure 10: Evaluation of MARL methods on alternate laydowns for *1o_10b_vs_1r*.

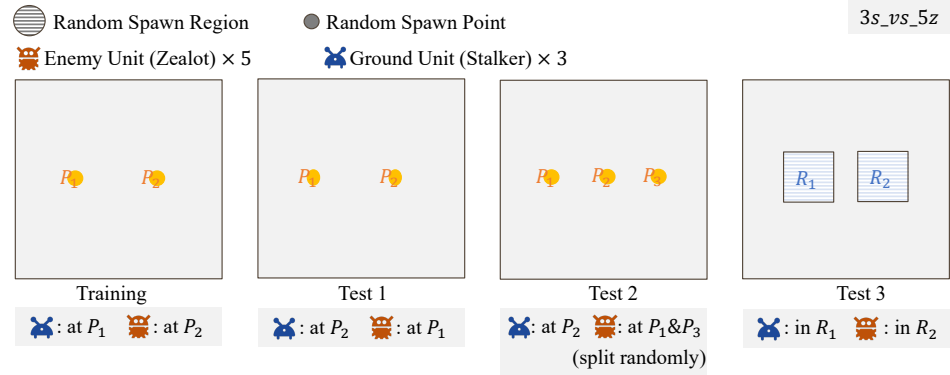
A.3 Visualization of OOD Laydowns



(a) 1o_10b_vs_1r



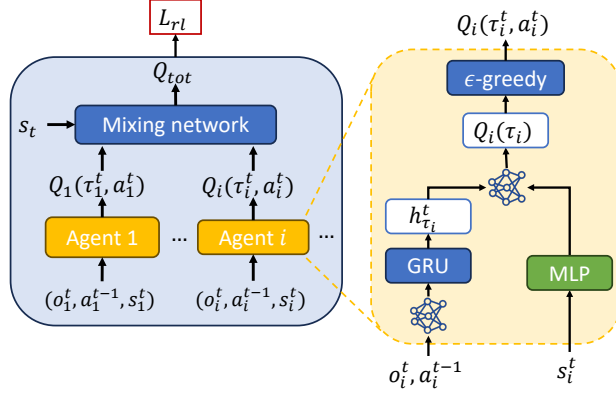
(b) 1o_2r_vs_4r



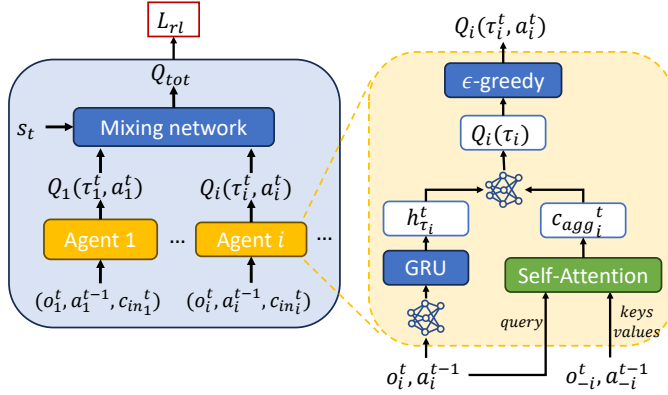
(c) 3s_vs_5z

Figure 11: Illustration of modified initial distributions for 1o_10b_vs_1r, 1o_2r_vs_4r and 3s_vs_5z. In this figure, "at" means that all indicated agents are generated at the same point P , while "in" denotes that all indicated agents are randomly generated at distinct locations within the region R .

A.4 Pipelines of QMIX-based Baselines



(a) QMIX-wState: the agent takes s_i^t as extra input at each time step.



(b) QMIX-Att: the agents receives all other agents' observation and action as communication messages.

Figure 12: Pipelines of QMIX-based baseline algorithms.

A.5 Evaluate Message Dropping

To assess the significance of communication, we conducted experiments in which all messages were disabled during test episodes. We observe from $3s_vs_5z$ test 2 that the performance drops approximately 20% in terms of win rate without communication. In other maps, the enemy is mostly stationary, so the information is primarily conveyed through the initial state and our approach effectively integrates this initial state information. However, even in such scenarios, we observed that the performance with communication was either equal to or superior to the performance without communication. This suggests that communication offers benefits beyond the initial state information, enhancing overall system performance.




Figure 13: Evaluation results of our algorithm with and without communication.

A.6 QMIX with Varying Observation Sight Ranges

Previous works [6, 10] discussed the dilemma in sight ranges: agents with small sight ranges can only observe limited information, making it challenging for them to engage in effective cooperation with their teammates. Conversely, agents with extensive sight ranges are more susceptible to distractions, degrading the cooperation quality. However, they only draw conclusions from the traffic junction and collaborative resource collection environments.

We conduct similar experiments on the *1o_10b_vs_1r* map by running QMIX with varying observation sight ranges. Interestingly, we have observed that a larger sight range results in quicker convergence, an increased win rate, and improved generalization to out-of-distribution (OOD) scenarios. We posit that this phenomenon is intertwined with the characteristics of the SMAC environment. In this context, effective cooperation among allies is critical for locating and attacking enemies to secure victory, and having a broader perspective on the overall situation is essential for accomplishing these complex tasks.

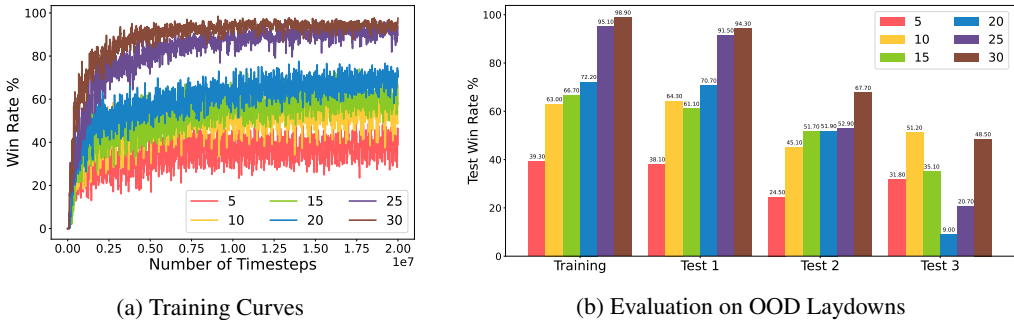


Figure 14: Training and evaluation results of QMIX with different sight ranges on *1o_10b_vs_1r*. A larger sight range gives us faster convergence speed and a higher win rate, which indicates that the agent can make wiser decisions if it is more aware of the global state under SMAC scenarios.

Sight Range		Training	Test 1	Test 2	Test 3
5	Win Rate	39.30%	38.10%	24.50%	31.80%
	Return	9.74 ± 8.77	9.52 ± 8.73	7.45 ± 7.61	8.65 ± 8.30
10	Win Rate	63.00%	64.30%	45.10%	51.20%
	Return	14.12 ± 8.42	14.40 ± 8.30	11.13 ± 8.62	12.34 ± 8.51
15	Win Rate	66.70%	61.10%	51.70%	35.10%
	Return	15.50 ± 7.17	14.73 ± 7.30	13.38 ± 7.46	9.86 ± 8.02
20	Win Rate	72.20%	70.70%	51.90%	9.00%
	Return	16.59 ± 6.28	16.37 ± 6.40	13.78 ± 7.00	6.53 ± 4.77
25	Win Rate	95.10%	91.50%	52.90%	20.70%
	Return	19.74 ± 2.84	19.26 ± 3.66	13.86 ± 6.95	8.55 ± 6.32
30	Win Rate	98.90%	94.30%	67.70%	48.50%
	Return	20.21 ± 1.32	19.62 ± 3.01	16.04 ± 6.25	13.28 ± 6.99

Table 1: Evaluations results on *1o_10b_vs_1r* over 20,000 episodes.

A.7 Additional Evaluation Details

Embedding Size		Training	Test 1	Test 2	Test 3
4	Win Rate	99.20%	95.40%	62.30%	63.30%
	Return	20.26 ± 1.14	19.77 ± 2.73	14.85 ± 7.22	15.21 ± 6.92
8	Win Rate	97.50%	96.20%	66.30%	69.10%
	Return	20.04 ± 1.98	19.87 ± 2.43	15.48 ± 6.93	16.25 ± 6.23
16	Win Rate	97.00%	96.20%	60.90%	73.30%
	Return	19.98 ± 2.19	19.89 ± 2.44	14.54 ± 7.40	16.74 ± 6.12
32	Win Rate	98.10%	96.70%	70.00%	70.70%
	Return	20.10 ± 1.75	19.93 ± 2.27	16.11 ± 6.56	16.47 ± 6.14

Table 2: Comparison of our algorithm with different embedding sizes on *1o_10b_vs_1r* over 20,000 episodes.

		Training	Test 1	Test 2	Test 3
QMIX	Win Rate	62.60%	60.40%	37.90%	27.80%
	Return	13.88 ± 8.64	13.59 ± 8.63	9.78 ± 8.52	8.73 ± 7.62
QMIX-Attn	Win Rate	94.80%	73.80%	35.20%	43.20%
	Return	19.69 ± 2.93	16.23 ± 7.10	9.63 ± 8.15	11.19 ± 8.28
QMIX-wState	Win Rate	96.60%	94.20%	55.80%	60.20%
	Return	19.92 ± 2.39	19.60 ± 3.12	13.31 ± 8.11	14.61 ± 7.27
MASIA	Win Rate	86.16%	68.93%	36.15%	16.46%
	Return	18.56 ± 4.56	15.59 ± 7.34	9.86 ± 8.16	6.7798 ± 6.50
TarMAC	Win Rate	90.95%	78.31%	31.13%	18.57%
	Return	19.25 ± 3.81	17.14 ± 6.42	8.93 ± 7.97	6.63 ± 7.00
NDQ	Win Rate	88.94%	68.54%	19.29%	0.32%
	Return	18.93 ± 4.18	15.74 ± 7.01	6.96 ± 6.99	1.68 ± 2.22
Our Method	Win Rate	98.00%	95.60%	68.60%	71.40%
	Return	20.11 ± 1.77	19.81 ± 2.63	15.79 ± 6.90	16.42 ± 6.37

Table 3: Evaluation results of various algorithms on *1o_10b_vs_1r* over 20,000 episodes.

		Training	Test 1	Test 2	Test 3
QMIX	Win Rate	50.60%	46.50%	23.00%	12.10%
	Return	10.75 ± 9.58	9.86 ± 9.65	7.74 ± 7.46	7.80 ± 5.15
QMIX-Attn	Win Rate	85.60%	46.20%	18.40%	8.40%
	Return	18.27 ± 4.28	9.81 ± 9.64	7.09 ± 6.91	7.10 ± 4.68
QMIX-wState	Win Rate	87.30%	78.20%	49.30%	15.50%
	Return	18.51 ± 3.96	16.82 ± 6.28	13.18 ± 7.01	8.78 ± 5.30
MASIA	Win Rate	82.81%	45.41%	22.21%	12.53%
	Return	17.93 ± 4.64	9.51 ± 9.73	7.19 ± 7.51	7.39 ± 5.43
TarMAC	Win Rate	76.28%	49.20%	20.82%	4.02%
	Return	16.60 ± 6.34	10.56 ± 9.53	7.11 ± 7.29	5.79 ± 3.96
NDQ	Win Rate	79.86%	43.20%	22.48%	8.95%
	Return	17.56 ± 4.97	9.28 ± 9.57	7.36 ± 7.46	6.72 ± 4.95
Our Method	Win Rate	87.60%	78.00%	61.20%	19.90%
	Return	18.61 ± 3.75	16.71 ± 6.50	14.92 ± 6.70	9.61 ± 5.54

Table 4: Evaluation results of various algorithms on *1o_2r_vs_4r* over 20,000 episodes.

		Training	Test 1	Test 2	Test 3
QMIX	Win Rate	99.80%	63.00%	94.40%	88.90%
	Return	21.17 ± 0.62	20.03 ± 4.05	21.66 ± 1.54	20.86 ± 2.50
QMIX-Attn	Win Rate	99.90%	0.00%	97.30%	80.00%
	Return	22.62 ± 0.44	0.01 ± 0.23	22.43 ± 1.30	20.94 ± 3.67
QMIX-wState	Win Rate	100.00%	16.20%	74.80%	51.80%
	Return	25.26 ± 1.13	18.09 ± 7.63	24.89 ± 3.57	17.39 ± 9.64
MASIA	Win Rate	0.00%	0.00%	0.00%	0.01%
	Return	9.81 ± 0.39	5.91 ± 2.08	5.28 ± 1.74	7.73 ± 1.41
TarMAC	Win Rate	94.59%	0.00%	24.37%	51.22%
	Return	21.67 ± 1.29	0.00 ± 0.11	18.84 ± 4.14	19.30 ± 4.60
NDQ	Win Rate	94.25%	0.13%	80.53%	19.37%
	Return	22.08 ± 1.50	6.10 ± 2.59	21.72 ± 2.89	14.99 ± 5.23
Our Method	Win Rate	99.70%	12.90%	96.00%	80.50%
	Return	20.95 ± 0.60	12.37 ± 7.28	21.22 ± 1.35	20.52 ± 3.26

Table 5: Evaluation results of various algorithms on *3s_vs_5z* over 20,000 episodes.

A.8 Hallucination

We also investigate the occurrence of false positives (hallucinations) at different time steps in the "1o_10b_vs_1r" scenario. By employing the disentangle loss to enforce local consistency, we have observed a noteworthy outcome: no hallucinations transpire within the agents' field of view. The instances of hallucinations that do occur outside the field of view can be primarily attributed to agents that are no longer active or are considered "dead." These hallucinations tend to persist for a few time steps before fading away.




Figure 15: False positives across timesteps. $< v$ means within the field of view, and $> v$ means outside the field of view.