

Revisiting Communication Efficiency in Multi-Agent Reinforcement Learning from the Dimensional Analysis Perspective

Chuxiong Sun*

National Key Laboratory of Space Integrated Information System, Institute of Software Chinese Academy of Sciences
Beijing, China
chuxiong2016@iscas.ac.cn

Peng He*

Beijing University of Posts and Telecommunications
Beijing, China
hepeng123@bupt.edu.cn

Rui Wang†

National Key Laboratory of Space Integrated Information System, Institute of Software Chinese Academy of Sciences
State Key Laboratory of Intelligent Game
Beijing, China
wangrui@iscas.ac.cn

Changwen Zheng

National Key Laboratory of Space Integrated Information System, Institute of Software Chinese Academy of Sciences
Beijing, China
changwen@iscas.ac.cn

ABSTRACT

In this work, we introduce a novel perspective—dimensional analysis—to address the challenge of communication efficiency in Multi-Agent Reinforcement Learning (MARL). Our findings reveal that simply optimizing the content and timing of communication at sending end is insufficient to fully resolve communication efficiency issues. Even after applying optimized and gated messages, dimensional redundancy and confounders still persist in the integrated message embeddings at receiving end, which negatively impact communication quality and decision-making. To address these challenges, we propose Dimensional Rational Multi-Agent Communication (DRMAC), designed to mitigate both dimensional redundancy and confounders in MARL. DRMAC incorporates a redundancy-reduction regularization term to encourage the decoupling of information across dimensions within the learned representations of integrated messages. Additionally, we introduce a dimensional mask that dynamically adjusts gradient weights during training to eliminate the influence of decision-irrelevant dimensions. We evaluate DRMAC across a diverse set of multi-agent tasks, demonstrating its superior performance over existing state-of-the-art methods in complex scenarios. Furthermore, the plug-and-play nature of DRMAC’s key modules highlights its generalizable performance, serving as a valuable complement rather than a replacement for existing multi-agent communication strategies.

KEYWORDS

Multi-Agent Reinforcement Learning; Multi-Agent Communication; Cooperative Multi-Agent Systems

*Equal contribution.

†Corresponding author.



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

ACM Reference Format:

Chuxiong Sun, Peng He, Rui Wang, and Changwen Zheng. 2025. Revisiting Communication Efficiency in Multi-Agent Reinforcement Learning from the Dimensional Analysis Perspective. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025)*, Detroit, Michigan, USA, May 19 – 23, 2025, IFAAMAS, 10 pages.

1 INTRODUCTION

Reinforcement Learning (RL) has achieved significant milestones across a wide array of complex real-world domains, including Game AI [25, 30, 31, 42], Robotics [1], and Autonomous Driving [18]. However, in cooperative multi-agent settings, distinct challenges emerge. One of the most prominent is the issue of partial observability, where agents are restricted to their local observations, lacking a comprehensive view of the entire environment. Further complicating the situation, Multi-Agent Reinforcement Learning (MARL) faces the problem of non-stationarity, as the environment’s dynamics constantly shift from the perspective of individual agents, thereby introducing significant complexity into the learning process. Multi-agent communication offers a promising solution to these challenges by enabling agents to gain a more comprehensive understanding of their surroundings through shared information. This approach not only stabilizes the learning process but also fosters coordinated actions among agents, ultimately improving overall system performance.

An essential challenge in multi-agent communication is optimizing communication efficiency. A critical observation is that only a subset of agents provide meaningful insights, while overwhelming the network with redundant information can be not only ineffective but also detrimental to learning [14]. Consequently, existing approaches have focused on refining both the **content** and **timing** of communication [15, 32, 34, 45, 51, 54]. These methods typically employ regularization techniques to reduce message redundancy or assess the relevance of observed information to determine optimal timing and partner selection for communication, thus minimizing unnecessary exchanges. We categorize these approaches as message-level de-redundancy strategies.




Figure 1: Existing communication methods typically focus on Step 1 and Step 2, optimizing the content, timing, and selection of partners for communication, with the aim of enhancing efficiency at the message level. However, in this work, we observe that even after these optimizations, dimensional redundancy and confounders still exist in the integrated message embeddings at Step 3. To further improve communication efficiency, we introduce a novel perspective—dimensional analysis—as a complement to existing communication methods.

Although message-level de-redundancy strategies have made significant strides, state-of-the-art methods continue to face challenges related to communication inefficiency. Hence, in this work, we revisit the multi-agent communication paradigm. As illustrated in Fig.1, once agents optimize the content and timing of their communication and share the messages with teammates, the receiving agents employ an encoder to generate embeddings that assimilate these messages. The essence of multi-agent communication lies in utilizing these embeddings to enhance agents' understanding of an uncertain environment, ultimately leading to more refined decision-making policies. Hence, in this work, we conduct dimensional analysis on these embeddings, identifying **dimensional redundancy** and **dimensional confounders** as the fundamental issues that hinder efficient communication.

Specifically, from an information theory perspective, each dimension encapsulates a portion of the representation's information entropy. Dimensional redundancy occurs when multiple dimensions contain overlapping information entropy. To demonstrate the existence of such dimensional redundancy, we visualize the integrated representations produced by both MASIA[11] and SMS[48]. MASIA represents a full-communication approach, while SMS employs a gating mechanism to filter messages at sending end. As illustrated in Fig.2, despite messages being gated and optimized at the sending end, substantial dimensional redundancy still persists at the receiving end. Therefore, it remains essential for the receiving agents to perform global filtering of all received information to further enhance communication efficiency.

The dimensional confounder refers to a set of dimensions containing "harmful" information, which ultimately degrades the model's performance. To demonstrate the existence of dimensional confounders, we conduct motivating experiments using both MASIA

and SMS. Specifically, we randomly mask several dimensions of message representation at receiving end by setting their values to zero, and subsequently evaluate the performance of the masked message representations. The experimental results are presented in Fig.3, where each point represents the outcome for several evaluated trajectories at a given mask ratio. Interestingly, representations of received messages with certain masked dimensions can achieve better performance compared to their unmasked counterparts. This empirical evidence supports our hypothesis regarding dimensional confounders, indicating that even after agents have filtered information during transmission, some decision-irrelevant information remains embedded in the representations integrated by the receiving agents.

With this vision in mind, we propose Dimensional Rational Multi-Agent Communication (DRMAC) to address both dimensional redundancy and confounders in MARL. Specifically, we propose a dimensional redundancy reduction regularization term that drives the cross-correlation matrix of twin message embeddings towards an identity matrix. The twin message embeddings are generated by taking the received messages as input along with a perturbed encoder of the same input, thereby creating two correlated views for contrastive learning. This simple yet powerful approach effectively reduces dimensional redundancy in the integrated message embeddings. To further combat the influence of dimensional confounders, we introduce a dimensional mask, which enhances the gradient effect of decision-relevant dimensions while reducing the influence of confounding dimensions during training. The learnable dimensional mask is optimized through a meta-learning paradigm, with the goal of enhancing the performance of masked representations on MARL tasks. Consequently, the dimensional mask continuously adjusts the weights of various dimensions based on their gradient contributions, encouraging the model to focus on acquiring specific decision-relevant information. This approach effectively addresses communication inefficiencies, particularly in complex tasks involving numerous agents and communication messages. We rigorously evaluated DRMAC across multiple MARL environments, including Hallway and SMAC. When compared to leading multi-agent communication strategies such as TarMac [7], MAIC [51], SMS [48], and MASIA [11], DRMAC consistently demonstrated superior performance and efficiency, highlighting its robustness and effectiveness in diverse settings. Our contributions can be summarized as follows:

- We revisit the existing paradigm of multi-agent communication and identify that optimizing the timing and content of messages at the receiving end alone is insufficient to address communication efficiency issues. Therefore, we introduce a novel perspective—dimensional analysis—to enhance the performance of multi-agent communication in MARL.
- We propose DRMAC, which incorporates a redundancy reduction regularization term and an Information Selective Network (ISN) to address the challenges of dimensional redundancy and dimensional confounders, respectively.
- Empirically, we demonstrate that DRMAC not only facilitates more effective message integration but also significantly enhances communication efficiency, thereby addressing a critical research gap in MARL.




Figure 2: The visualizations depict the representations learned by SMS[48], MASIA[11] and our method on a challenging SMAC task, 1o_2r_vs_4r. The learned features are projected into an RGB color image, where distinct colors indicate different feature types. The horizontal axis corresponds to feature dimensions, while the vertical axis represents samples from different trajectories. Greater color contrast signifies lower similarity between feature dimensions. These plots illustrate the similarity between dimensional features within a batch. In contrast to existing approaches, our method employs a redundancy-reduction technique that efficiently decouples the received information into distinct dimensions. Each of these dimensions represents a unique part of the information’s entropy, ensuring that the resulting representation is both more informative and less redundant.




Figure 3: Experimental scatter plots were generated by MASIA and SMS with randomly masked dimensions on a challenging SMAC task, 1o_2r_vs_4r. The Baseline and the red dashed lines indicate the performance achieved by the unmasked representation of MASIA and SMS. Each point represents an independent experimental result, obtained by applying a specific mask rate to the original representation at the dimensional level. Notably, the original representation remains unchanged throughout the experiments. These results demonstrate the pervasive presence of dimensional confounders in the process of multi-agent communication.

2 RELATED WORKS

2.1 Multi-Agent Communication

Multi-agent communication has emerged as an indispensable component in MARL[22, 27, 38, 44, 50]. Research in this domain has primarily concentrated on three fundamental questions:

What to communicate. CommNet [34], as a pioneering work in this domain, enabled agents to learn continuous messages for effective communication. Building on the foundation laid by CommNet, numerous approaches have been proposed to further refine the message learning process. VBC [53] focuses on filtering out noise while preserving valuable information by restricting message variance. TMC [54] introduces regularizers to minimize temporal redundancy in messages, thereby enhancing efficiency. NDQ [45] leverages information-theoretic regularizers to generate expressive

yet concise messages. MAIC [51] advances communication by allowing agents to customize messages for specific recipients, thereby promoting tailored and more effective information exchange.

When and whom to communicate. To improve communication efficiency, methods like IC3Net [32] and ATOC [15] have employed gating networks to eliminate unnecessary communication links. Similarly, SchedNet [15] and IMMAC [36] utilize heuristic mechanisms to evaluate the significance of observations, thereby gating non-essential communication and reducing overhead. Additionally, techniques such as MAGIC [23], I2C [8], SMS [48] and T2MAC[37] have been developed to identify the most appropriate recipients for messages. These methods focus on assessing the impact of shared information on the recipients’ decision-making, ensuring that communication is directed where it can most effectively enhance decision-making outcomes.

How to utilize received information. Existing works [7, 46] have explored how agents can effectively assimilate crucial information

from an abundance of raw messages. MASIA [11] and M2I2[35] represent the pioneering effort to introduce self-supervised representation learning into multi-agent communication. Specifically, these works proposed self-supervised auxiliary tasks to enhance the efficiency of information integration. However, this paper primarily focuses on addressing the issues of dimensional redundancy and dimensional confounders in multi-agent communication representation learning.

2.2 Self-supervised learning

Self-supervised learning approaches [6, 9, 10, 13, 19, 20, 28, 40, 41, 43, 47, 52] have achieved significant breakthroughs in unsupervised representation learning. Specifically, CMC [39] and AMDIM [2] leverage contrastive learning techniques on multi-view datasets, thereby enhancing representation quality by integrating multiple views. SimCLR [5] and MoCo [12] utilize large batch sizes or memory banks to expand the pool of negative samples, facilitating the learning of more discriminative and robust representations. SwAV [4], on the other hand, compares cluster assignments across different views instead of directly comparing features, thereby leveraging multiple views to improve learning efficiency. Methods such as BYOL [10] and Barlow Twins [52] have also identified a key challenge—insufficient self-supervision can lead to feature collapse, where learned representations degenerate into trivial solutions. Nonetheless, these benchmark self-supervised learning frameworks benefit from well-designed architectures that incorporate sufficient self-supervision, resulting in improved discriminative capabilities of the learned representations. Recently, these advantages have been extended to Visual Reinforcement Learning (VRL)[16, 17, 26, 49], where contrastive learning is employed to derive informative and discriminative low-dimensional representations from high-dimensional visual inputs. Additionally, MA2CL[33], a MARL method, has adopted self-supervised learning to encourage the learned representations to be both temporally and agent-level predictive by reconstructing masked agent observations within the latent space. Consequently, a growing number of RL approaches are integrating self-supervised learning techniques to develop more refined state representations. Unlike prior methods that primarily aim to enhance state representations by combining self-supervised learning with RL, DRMAC models the information integration challenge in multi-agent communication as a representation learning problem. Specifically, DRMAC seeks to enhance communication efficiency by reducing redundancy and mitigating confounders in the integrated representations, thereby significantly improving the quality of communication among agents.

3 BACKGROUND

In this work, we focus on fully cooperative multi-agent tasks, characterized by partial observability and necessity for inter-agent communication. These tasks are modeled as Decentralized Partially Observable Markov Decision Processes (Dec-POMDPs) [24], represented by the tuple $G = (N, S, O, A, \mathbb{O}, P, R, \gamma, M)$. In this formulation, $N \equiv \{1, \dots, n\}$ denotes the set of agents, S represents the global states, O describes the observations available to each agent, A signifies the set of available actions, $\mathbb{O}$ is the observation function

mapping states to observations, P is the transition function illustrating the dynamics of the environment, R is the reward function dependent on the global states and joint actions of the agents, γ is the discount factor, and M specifies the set of messages that can be communicated among the agents. At each time-step t , each agent $i \in N$ has access to its own observation $o_i^t \in O$ determined by the observation function $\mathbb{O}(o_i^t | s_t)$. Additionally, each agent can receive messages $c_i^t = \sum_{j \neq i} m_j^t$ from teammates $j \in N$. Utilizing both the observed and received information, agents then make local decisions. As each agent selects an action, the joint action a_t results in a shared reward $r_t = R(s_t, a_t)$ and transitions the system to the next state s_{t+1} according to the transition function $P(s_{t+1} | s_t, a_t)$. The objective for all agents is to collaboratively develop a joint policy π to maximize the discounted cumulative return $\sum_{t=0}^T \gamma^t r_t$.

4 METHODS

As illustrated in Fig. 4, upon receiving messages $(m_1^t, \dots, m_n^t)$, DRMAC incorporates an Information Integration Encoder (IIE) to aggregate the information shared by agents into a common representation z_t . Our goal is to reduce dimensional redundancy and confounders in this representation, thereby ensuring that each agent receives sufficient and relevant information to make effective decisions. The key features of DRMAC are highlighted as follows:

- **Redundancy Reduction Regularization:** DRMAC’s integration network at the receiving end is optimized using a redundancy reduction regularization term, which aims to decouple the received information into distinct dimensions, with each dimension encapsulating a unique subset of the representation’s information entropy.
- **Information Selective Network (ISN):** After encoding a decoupled representation, the ISN learns a dimensional mask that adjusts the gradient weights of each dimension according to their contributions to the decision-making process. This allows the policy network to focus on relevant and significant dimensional information while effectively ignoring dimensional confounders.

We provide the details of the redundancy reduction mechanism in **Section 4.1**, and the specifics of ISN in **Section 4.2**. Additionally, a comprehensive description of the training process is outlined in **Section 4.3**.

4.1 Dimensional Redundancy Reduction for Integrated Message Embeddings

As illustrated in Fig. 2, despite optimizing communication content and timing at message level, multi-agent communication methods still face the issue of dimensional redundancy—where multiple dimensions in the integrated message embeddings contain overlapping information entropy. To address this challenge, we draw inspiration from a neuroscience approach, specifically redundancy reduction [3], which hypothesizes that the goal of sensory processing is to recode highly redundant sensory inputs into a factorial code, consisting of statistically independent components. Building on its successful application in self-supervised learning [52], we apply dimensional redundancy reduction to integrated message




Figure 4: The overview of our proposed DRMAC.

embeddings to mitigate overlapping information and improve communication efficiency. Concretely, for agent i , the received messages are first fed into the Information Integration Encoder to obtain its representation z_i^t . To obtain multiple views of the message embeddings, we duplicate the parameters of the IIE and add Gaussian noise to the copied version. We then encode the identical set of received messages using both the original IIE and the noise-added IIE, resulting in two embeddings, z_i^t and $z_i^{t'}$. Essentially, this process serves as a form of data augmentation. Furthermore, previous works [5] have experimentally demonstrated that introducing a nonlinear layer before the contrastive loss can significantly improve performance. Therefore, we also include a non-linear projection layer. This layer is implemented as an MLP network. To simplify notations, we utilize z^A and z^B to denote the two embeddings after online projector and target projector respectively. The redundancy reduction objective is formulated as follows:

$$\mathcal{L}_{RR} \triangleq \underbrace{\sum_i (1 - C_{ii})^2}_{\text{invariance term}} + \underbrace{\lambda \sum_i \sum_{j \neq i} C_{ij}^2}_{\text{redundancy reduction term}} \quad (1)$$

where λ is a hyper-parameter, and C is computed between the outputs of the online and target projectors:

$$C_{ij} \triangleq \frac{\sum_b z_{b,i}^A z_{b,j}^B}{\sqrt{\sum_b (z_{b,i}^A)^2} \sqrt{\sum_b (z_{b,j}^B)^2}} \quad (2)$$

where b indexes batch samples, and i, j index the vector dimensions of the network outputs. The matrix C is a square matrix with a size equal to the dimensionality of the network's output, and its values range between -1 (indicating perfect anti-correlation) and 1 (indicating perfect correlation).

Intuitively, the invariance term of the objective aims to set the diagonal elements of the cross-correlation matrix to 1, thereby ensuring that the embeddings remain invariant to applied distortions. Meanwhile, the redundancy reduction term seeks to bring the off-diagonal elements of the cross-correlation matrix to 0, effectively decorrelating the different components of the embedding vector. This decorrelation minimizes redundancy among the output units, ensuring that each unit captures unique, non-overlapping information in the embeddings of the received messages. With the centralized process enhanced by $\mathcal{L}_{RR}$, the online Information Integration Encoder can efficiently integrate incoming messages and produce decoupled representations. It is important to note that this process is only applied during the training phase and is not utilized during execution.

4.2 Information Selective Network: Addressing Dimensional Confounders

To address the issue of dimensional confounders, we use the Information Selective Network (ISN) to construct a learnable dimensional mask $\omega = \{\omega_k \mid k \in \llbracket 1, D \rrbracket\}$. Considering that, in the CTDE framework, receiving agents operate in different situations and fulfill distinct roles within the team, the dimensional confounders in z may vary for each agent. Therefore, we design the ISN to take each agent's observation as input and produce a dimensional mask specific to that agent. The ISN is implemented as a Multi-Layer Perceptron (MLP) with a sigmoid activation function at the output layer, ensuring that each dimension of ω_k is bounded between 0 and 1. The dimensional mask is then used to assign a weight to each dimension of the representation.

$$\tilde{z}_i^t = z_i^t \otimes \omega_i^t \quad (3)$$

where $\tilde{z}_i^t$ denotes the masked representation, ω_i^t represents the dimensional mask for agent i at time-step t and $\otimes$ is an element-wise Hadamard product function. By applying element-wise multiplication to $\tilde{z}_i^t$, each agent derives a distinct representation tailored to its specific needs. This process enables each agent to extract only the information from the decoupled representation that is most relevant to its decision-making context, while ignoring dimensions that serve as confounders. Then, the masked representation, representing receiving agent's understanding about received messages are concatenated with its own local observations to make the local decisions, $a_i^t = \pi_{\theta_\pi}(o_i^t, \tilde{z}_i^t)$. This tailored approach ensures that each agent's representation is aligned with its unique requirements, thereby enhancing the quality of decision-making and improving overall efficiency.

4.3 Overall Training Process

The ISN is designed to explore dimensional confounders within multiple observed information, specifically tailoring the extracted representation to meet the diverse needs of local decision-making for multiple agents. The primary challenge here lies in the dynamic nature of the MARL environment and the variability in communication needs at different stages of a mission. In contrast to static environments, the significance of certain information dimensions can change dramatically, requiring ISN to continuously adapt and maintain its effectiveness. This complexity makes it difficult to solve the problem through traditional optimization methods, which often rely on first-order gradients and are insufficient in scenarios characterized by rapidly changing conditions and diverse agent requirements. Hence, we utilize meta-learning to update the ISN. Meta-learning provides the ISN with the ability to learn how to adapt quickly to new situations by leveraging past experiences. This approach enables the ISN to continuously adjust to the dynamic communication needs of agents, efficiently dealing with the ever-changing importance of different information dimensions.

In our training framework, it is crucial to emphasize that only the parameters of θ_{ISN} are refined through the meta-learning process, while the rest of the system parameters are updated via standard first-order gradient methods. Specifically, in the first regular training step, we concentrate on optimizing the parameter set $\theta = (\theta_{OE}, \theta_{IEE}, \theta_P, \theta_\pi)$ by jointly minimizing both the redundancy reduction loss and the RL loss, as described by:

$$\arg \min_{\theta} \mathcal{L}_{tot}(\theta, \theta_{ISN}), \quad (4)$$

where $\mathcal{L}_{tot}(\theta, \theta_{ISN}) = \mathcal{L}_{RL} + \beta \mathcal{L}_{RR}$, and β is a weighting coefficient to balance the RL objective with the redundancy reduction objective. It is worth noting that our DRMAC framework is compatible with any MARL training algorithm and multi-agent communication strategy. The experimental results presented in **Section 5.4** substantiate this assertion.

In the second meta-learning-based step, θ_{ISN} is updated using a meta-learning approach that leverages second-order gradient information [21]. This method is crucial for adapting θ_{ISN} to properly recognize the importance of each information dimension. The

update involves calculating the gradient of θ_{ISN} based on the combined performance metric $\mathcal{L}_{RL}$, and is formulated as follows:

$$\arg \min_{\theta_{ISN}} \mathcal{L}_{RL}(\theta_{trial}, \theta_{ISN}), \quad (5)$$

where $\theta_{trial} = (\theta_{OE}^{trial}, \theta_{IEE}^{trial}, \theta_P^{trial}, \theta_\pi^{trial})$ are the trial weights obtained after a single gradient update on $\mathcal{L}_{RL}$. This update of trial weights is computed as:

$$\theta_{trial} = \theta - \ell_\theta \nabla_{\theta} \mathcal{L}_{RL}, \quad (6)$$

with ℓ_θ representing the learning rate. Notably, during the calculation of these trial weights, back-propagation of gradients is excluded to ensure computational efficiency. Consequently, θ_{ISN} is refined through second-order gradient optimization based on θ . This process allows ISN to be continuously fine-tuned by considering the gradient contributions from $\mathcal{L}_{RL}$, enabling it to effectively discern the relevance of each information dimension and adapt accordingly. By using this meta-learning-based approach, ISN dynamically identifies and focuses on critical dimensions, enhancing decision-making quality and ensuring communication efficiency.

We train the dimensional mask ω solely based on RL performance for two key reasons. First, dimensional redundancy reduction acts as a regularization term that imposes strong constraints. Second, redundancy reduction is independent of downstream decision-making tasks, while our primary goal is to use the masked representation to improve agents' understanding of environmental uncertainty and enhance cooperative decision-making. By focusing on RL performance to train ω , we enable agents to effectively identify and mitigate dimensional confounders that are irrelevant to cooperative decision-making.

5 EXPERIMENTS

In this section, we present our experimental design, which is meticulously structured to address three fundamental research questions:

- **RQ1.** How does the performance of DRMAC compare to that of state-of-the-art communication methods?
- **RQ2.** Which specific components of DRMAC are crucial to its overall performance?
- **RQ3.** Can DRMAC be effectively integrated with various existing communication baselines to enhance their communication efficiency?

5.1 Setup

Baselines. To address RQ1, we compare MASIA against a diverse set of baselines, including a strong communication-free baseline, QMIX [27], as well as several state-of-the-art communication methods, such as MASIA [11], MAIC [51], SMS [48], and TarMAC [7]. In our experiments, the implementation of each baseline method strictly follows the official repositories, with all hyperparameters left at their original, best-performing configurations to ensure fair and consistent comparison.

Environments. We conduct a comprehensive evaluation of DRMAC across two well-known cooperative multi-agent environments, encompassing a total of eight tasks that vary in the number of agents, difficulty levels, and task types. We begin with the Hallway environment [45], which provides a relatively simple setup based on multiple Markov chains. In this environment, agents are




Figure 5: Performance on multiple tasks.

randomly positioned within different chains and must collaborate to reach a goal state simultaneously, operating under conditions of partial observability. To further assess cooperative performance, we introduce an improved variant named HallwayGroup, featuring diverse agents grouped into different cohorts, each with distinct requirements for reaching their respective goal states at different times. SMAC [29] is derived from the well-known real-time strategy game StarCraft II, focusing on micromanagement challenges where each unit is independently controlled by an agent making decisions under partial observability. To increase task difficulty, we selected three "superhard" tasks, which are particularly challenging due to their complex opponent compositions and dynamic game states. Additionally, to further stress the difficulties inherent in multi-agent communication, we incorporated three more tasks from the setup described in [45]. These tasks not only restrict the agents' vision range but also immerse them in highly intricate maps, characterized by maze-like terrain or unpredictable spawning dynamics of both allied and enemy units. This combination of limited visibility and complex environments provides a rigorous assessment of DRMAC's ability to manage coordination under demanding conditions.

Implementation. Our DRMAC implementation is built upon PyMARL and QMIX, utilizing a basic full communication strategy at the sending end.

5.2 Performance(RQ1)

Our evaluation begins with a comparative analysis of the learning curves of DRMAC against several baseline methods across diverse environments. This comparison aims to assess DRMAC's overall cooperative performance and communication efficiency. As shown in Fig. 5, DRMAC consistently outperforms the baselines

in almost all environments, demonstrating its robust capabilities across a wide range of cooperative multi-agent tasks of varying difficulties and scales. This superior performance can be attributed to DRMAC's ability to effectively counteract dimensional redundancy and confounders, allowing it to consistently identify the most relevant dimensions in the integrated message embeddings at the receiving end. It is worth noting that, across nearly all environments, the baselines employing communication significantly outperform the communication-free baseline. This result suggests that the selected environments and tasks are inherently communication-intensive and challenging, further reinforcing DRMAC's superior performance in handling such complex scenarios. To gain a deeper understanding of the intrinsic mechanisms behind DRMAC's performance, we conduct a more detailed analysis of specific tasks. In SMAC, five of the six selected tasks, excluding 5z_vs_1ul, involve controlling diverse types of agents. For these tasks, existing methods stress the importance of designing communication architectures (e.g., SMS) and tailoring communication content (e.g., MAIC) to align with the specific capabilities of each unit. However, we observe that both MAIC and SMS do not perform consistently well across these five tasks. For instance, while MAIC shows strong performance in tasks like 3s5z_vs_3s6z and MMM2, it underperforms in other heterogeneous scenarios. In contrast, DRMAC, as well as MASIA, consistently demonstrates superior performance across all tasks. This observation further supports our analysis that optimizing communication filtering and content at the sending end is only the first step toward enhancing communication efficiency. Even after optimizing the communication content and timing, receiving agents must still carefully integrate the received information to form embeddings that are truly useful for




Figure 6: Performance of different DRMAC variants.

decision-making. Moreover, the consistently superior performance of DRMAC underscores the importance of addressing dimensional redundancy and confounders, highlighting DRMAC’s effectiveness in overcoming these challenges.

5.3 Ablation(RQ2)

We conduct several ablation studies, primarily focusing on the scenario *1o_2r_vs_4r*, to evaluate the specific contributions of different components in DRMAC. Below is a breakdown of the configurations examined:

- **DRMAC**: This refers to the complete method proposed in our work, incorporating all components.
- **QMIX**: Serving as a baseline, QMIX represents the core functionality without the enhancements introduced by DRMAC.
- **DRMAC w/o RR**: This is a variant of DRMAC where the redundancy reduction techniques are removed, allowing us to assess the impact of dimensional redundancy reduction.
- **DRMAC w/o ISN**: This variant of DRMAC omits ISN, treating each dimension of the representation as equally important. This allows us to assess the significance of mitigating dimensional confounders.

As illustrated in Figure 6, the performance of both DRMAC w/o RR and DRMAC w/o ISN is suboptimal. The poor performance of DRMAC w/o ISN aligns with our expectations, as redundancy reduction primarily aims to decouple the received information, which is not directly relevant to decision-making performance. However, the underperformance of DRMAC w/o RR further underscores the significance of the proposed redundancy reduction regularization. This result indicates that relying solely on ISN is insufficient to achieve state-of-the-art performance. We attribute this to the presence of dimensional redundancy, which forces ISN to identify agent-specific dimensional confounders from high-dimensional redundant representations, thereby increasing the complexity of training ISN. In summary, our approach of extracting dimensional confounders from redundancy-reduced representations is mutually reinforcing, with both components being essential for optimal performance.

5.4 Generalizability(RQ3)

We begin by evaluating the generalizability of DRMAC by combine it with different communication-free MARL methods, such as VDN[38], QMIX[27] and MAPPO[50]. As illustrated in Fig.7,




Figure 7: Performance of DRMAC integrated with various MARL baselines.

DRMAC consistently delivers superior performance across all these baselines, often with a significant improvement margin. This positive performance enhancement underscores the broad applicability and effectiveness of DRMAC in addressing dimensional redundancy and confounders, ultimately leading to improved communication efficiency and cooperative performance in MARL.

To further demonstrate the generalizability of DRMAC, we implement it as a plug-and-play module into several state-of-the-art multi-agent communication methods, specifically MASIA[11], SMS[48], and TarMAC[7]. Among these methods, MASIA and TarMAC are fully communication-based, whereas SMS incorporates message filtering at the sender side. As illustrated in Fig. 7, we observe that DRMAC, as a plug-and-play module, consistently enhances the communication performance of these state-of-the-art methods. This result further validates the presence of dimensional redundancy and confounders, as well as their negative impact on communication efficiency. Moreover, these findings highlight the significance and effectiveness of studying the dimensional rationale at the receiving end, regardless of whether message-level filtering has already been applied at the sender side. In other words, applying DRMAC to conduct dimensional analysis on the representations after the received messages are encoded can serve as a complementary, rather than a replacement, to existing approaches.

6 CONCLUSION

In this work, we addressed the challenge of communication efficiency in MARL by introducing a novel perspective—dimensional analysis—to complement existing methods. Our experimental results demonstrate that optimizing the content and timing of communication alone is insufficient to fully address the issue of communication efficiency, as dimensional redundancy and confounders continue to persist. To overcome these challenges, we proposed DRMAC, a method designed to mitigate both dimensional redundancy and confounders in MARL. Specifically, we introduced a redundancy-reduction regularization term that encourages the dimensions of the learned message representations to model decoupled information. Furthermore, we employed a dimensional mask to adjust the gradient weights of different dimensions during training, thereby effectively eliminating the influence of decision-irrelevant dimensions. Our approach not only enhances the efficiency of communication but also leads to more robust and effective decision-making in MARL.

REFERENCES

- [1] OpenAI: Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Jozefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, et al. 2020. Learning dexterous in-hand manipulation. *The International Journal of Robotics Research* 39, 1 (2020), 3–20.
- [2] Philip Bachman, R. Devon Hjelm, and William Buchwalter. 2019. Learning Representations by Maximizing Mutual Information Across Views. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8–14, 2019, Vancouver, BC, Canada*, Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett (Eds.). 15509–15519. <https://proceedings.neurips.cc/paper/2019/hash/ddf354219aac374f1d40b7e760ee5bb7-Abstract.html>
- [3] Horace B Barlow et al. 1961. Possible principles underlying the transformation of sensory messages. *Sensory communication* 1, 01 (1961), 217–233.
- [4] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. 2020. Unsupervised Learning of Visual Features by Contrasting Cluster Assignments. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6–12, 2020, virtual*, Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.). <https://proceedings.neurips.cc/paper/2020/hash/70feb62b69f16e0238f741fab228fec2-Abstract.html>
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. 2020. A Simple Framework for Contrastive Learning of Visual Representations. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13–18 July 2020, Virtual Event (Proceedings of Machine Learning Research, Vol. 119)*, PMLR, 1597–1607. <http://proceedings.mlr.press/v119/chen20j.html>
- [6] Ching-Yao Chuang, Joshua Robinson, Yen-Chen Lin, Antonio Torralba, and Stefanie Jegelka. 2020. Debaised Contrastive Learning. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6–12, 2020, virtual*, Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.). <https://proceedings.neurips.cc/paper/2020/hash/63c3ddcc7b23daa1e42dc419a44a873-Abstract.html>
- [7] Abhishek Das, Théophile Gervet, Joshua Romoff, Dhruv Batra, Devi Parikh, Mike Rabbat, and Joelle Pineau. 2019. Tarmac: Targeted multi-agent communication. In *International Conference on Machine Learning*. 1538–1546.
- [8] Ziluo Ding, Tiejun Huang, and Zongqing Lu. 2020. Learning individually inferred communication for multi-agent cooperation. *Advances in Neural Information Processing Systems* 33 (2020), 22069–22079.
- [9] Aleksandr Ermolov, Aliaksandr Siarohin, Enver Sangineto, and Nicu Sebe. 2021. Whitening for Self-Supervised Representation Learning. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18–24 July 2021, Virtual Event (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 3015–3024. <http://proceedings.mlr.press/v139/ermolov21a.html>
- [10] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Ávila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. 2020. Bootstrap Your Own Latent - A New Approach to Self-Supervised Learning. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6–12, 2020, virtual*, Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.). <https://proceedings.neurips.cc/paper/2020/hash/f3ada80d54ee70142b17b8192b2958e-Abstract.html>
- [11] Cong Guan, Feng Chen, Lei Yuan, Chenghe Wang, Hao Yin, Zongzhang Zhang, and Yang Yu. 2022. Efficient Multi-agent Communication via Self-supervised Information Aggregation. *Advances in Neural Information Processing Systems* 35 (2022), 1020–1033.
- [12] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. 2020. Momentum Contrast for Unsupervised Visual Representation Learning. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13–19, 2020*, Computer Vision Foundation / IEEE, 9726–9735. <https://doi.org/10.1109/CVPR42600.2020.00975>
- [13] Qirui Ji, Jiangmeng Li, Jie Hu, Rui Wang, Changwen Zheng, and Fanjiang Xu. 2024. Rethinking Dimensional Rationale in Graph Contrastive Learning from Causal Perspective. In *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2024, February 20–27, 2024, Vancouver, Canada*, Michael J. Wooldridge, Jennifer G. Dy, and Srirama Natarajan (Eds.). AAAI Press, 12810–12820. <https://doi.org/10.1609/AAAI.V38I1.29177>
- [14] Jiechuan Jiang and Zongqing Lu. 2018. Learning attentional communication for multi-agent cooperation. In *Advances in neural information processing systems*. 7254–7264.
- [15] Daewoo Kim, Sangwoo Moon, David Hostallero, Wan Ju Kang, Taeyoung Lee, Kyunghwan Son, and Yung Yi. 2019. Learning to schedule communication in multi-agent reinforcement learning. *arXiv preprint arXiv:1902.01554* (2019).
- [16] Michael Laskin, Kimin Lee, Adam Stooke, Lerrel Pinto, Pieter Abbeel, and Aravind Srinivas. 2020. Reinforcement Learning with Augmented Data. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6–12, 2020, virtual*, Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.). <https://proceedings.neurips.cc/paper/2020/hash/e615c82aba461681ade82da2da38004a-Abstract.html>
- [17] Michael Laskin, Aravind Srinivas, and Pieter Abbeel. 2020. CURL: Contrastive Unsupervised Representations for Reinforcement Learning. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13–18 July 2020, Virtual Event (Proceedings of Machine Learning Research, Vol. 119)*, PMLR, 5639–5650. <http://proceedings.mlr.press/v119/laskin20a.html>
- [18] Edouard Leurent. 2018. A survey of state-action representations for autonomous driving. (2018).
- [19] Jiangmeng Li, Wenwen Qiang, Yanan Zhang, Wenyi Mo, Changwen Zheng, Bing Su, and Hui Xiong. 2022. MetaMask: Revisiting Dimensional Confounder for Self-Supervised Learning. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, Sammi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (Eds.). http://papers.nips.cc/paper_files/paper/2022/hash/fb575ab4d882a4c734641155a5f30911-Abstract-Conference.html
- [20] Jiangmeng Li, Wenwen Qiang, Changwen Zheng, Bing Su, and Hui Xiong. 2022. MetaAug: Contrastive Learning via Meta Feature Augmentation. In *International Conference on Machine Learning, ICML 2022, 17–23 July 2022, Baltimore, Maryland, USA (Proceedings of Machine Learning Research, Vol. 162)*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato (Eds.). PMLR, 12964–12978. <https://proceedings.mlr.press/v162/li22r.html>
- [21] Shikun Liu, Andrew Davison, and Edward Johns. 2019. Self-supervised generalisation with meta auxiliary learning. *Advances in Neural Information Processing Systems* 32 (2019).
- [22] Ryan Lowe, Yi I Wu, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Advances in neural information processing systems*. 6379–6390.
- [23] Yaru Niu, Rohan R Paleja, and Matthew C Gombolay. 2021. Multi-Agent Graph-Attention Communication and Teaming. In AAMAS. 964–973.
- [24] Frans A Oliehoek, Christopher Amato, et al. 2016. *A concise introduction to decentralized POMDPs*. Vol. 1. Springer.
- [25] Ian Osband, Charles Blundell, Alexander Pritzel, and Benjamin Van Roy. 2016. Deep exploration via bootstrapped DQN. In *Advances in neural information processing systems*. 4026–4034.
- [26] Roberta Raileanu, Maxwell Goldstein, Denis Yarats, Ilya Kostrikov, and Rob Fergus. 2021. Automatic Data Augmentation for Generalization in Reinforcement Learning. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6–14, 2021, virtual*, Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (Eds.). 5402–5415. <https://proceedings.neurips.cc/paper/2021/hash/2b38c2d2f6a49b97f706ec9148ce48d86-Abstract.html>
- [27] Tabish Rashid, Mikayel Samvelyan, Christian Schroeder De Witt, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. 2018. QMIX: Monotonic value function factorisation for deep multi-agent reinforcement learning. *arXiv preprint arXiv:1803.11485* (2018).
- [28] Joshua David Robinson, Ching-Yao Chuang, Suvrit Sra, and Stefanie Jegelka. 2021. Contrastive Learning with Hard Negative Samples. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3–7, 2021*, OpenReview.net. <https://openreview.net/forum?id=CR1XOQ0UTh>
- [29] Mikayel Samvelyan, Tabish Rashid, Christian Schroeder de Witt, Gregory Farquhar, Nantas Nardelli, Tim GJ Rudner, Chia-Man Hung, Philip HS Torr, Jakob Foerster, and Shimon Whiteson. 2019. The starcraft multi-agent challenge. *arXiv preprint arXiv:1902.04043* (2019).
- [30] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharmash Kumar, Thore Graepel, et al. 2018. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science* 362, 6419 (2018), 1140–1144.
- [31] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. 2017. Mastering the game of go without human knowledge. *nature* 550, 7676 (2017), 354–359.
- [32] Amanpreet Singh, Tushar Jain, and Sainbayar Sukhbaatar. 2018. Learning when to communicate at scale in multiagent cooperative and competitive tasks. *arXiv preprint arXiv:1812.09755* (2018).
- [33] Haolin Song, Mingxiao Feng, Wengang Zhou, and Houqiang Li. 2023. Ma2cl: Masked attentive contrastive learning for multi-agent reinforcement learning. *arXiv preprint arXiv:2306.02006* (2023).
- [34] Sainbayar Sukhbaatar, Arthur Szlam, and Rob Fergus. 2016. Learning Multiagent Communication with Backpropagation. In *Proceedings of the 30th International Conference on Neural Information Processing Systems (Barcelona, Spain) (NIPS'16)*, Curran Associates Inc., Red Hook, NY, USA, 2252–2260.

- [35] Chuxiong Sun, Peng He, Qirui Ji, Zehua Zang, Jiangmeng Li, Rui Wang, and Wei Wang. 2024. M2I2: Learning Efficient Multi-Agent Communication via Masked State Modeling and Intention Inference. *arXiv preprint arXiv:2501.00312* (2024).
- [36] Chuxiong Sun, Bo Wu, Rui Wang, Xiaohui Hu, Xiaoya Yang, and Cong Cong. 2021. Intrinsic Motivated Multi-Agent Communication. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems* (Virtual Event, United Kingdom) (AAMAS '21). International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1668–1670.
- [37] Chuxiong Sun, Zehua Zang, Jiabao Li, Jiangmeng Li, Xiao Xu, Rui Wang, and Changwen Zheng. 2024. T2MAC: Targeted and Trusted Multi-Agent Communication through Selective Engagement and Evidence-Driven Integration. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 15154–15163.
- [38] Peter Sunehag, Guy Lever, Audrunas Gruslys, Wojciech Marian Czarnecki, Viničius Zambaldi, Max Jaderberg, Marc Lanctot, Nicolas Sonnerat, Joel Z Leibo, Karl Tuyls, et al. 2017. Value-decomposition networks for cooperative multi-agent learning. *arXiv preprint arXiv:1706.05296* (2017).
- [39] Yonglong Tian, Dilip Krishnan, and Phillip Isola. 2020. Contrastive Multiview Coding. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XI (Lecture Notes in Computer Science, Vol. 12356)*, Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (Eds.). Springer, 776–794. https://doi.org/10.1007/978-3-030-58621-8_45
- [40] Yao-Hung Hubert Tsai, Yue Wu, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2021. Self-supervised Learning from a Multi-view Perspective. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net. https://openreview.net/forum?id=bdp_8ltjwp
- [41] Vikas Verma, Thang Luong, Kenji Kawaguchi, Hieu Pham, and Quoc V. Le. 2021. Towards Domain-Agnostic Contrastive Learning. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 10530–10541. <http://proceedings.mlr.press/v139/verma21a.html>
- [42] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* 575, 7782 (2019), 350–354.
- [43] Jingyao Wang, Wenwen Qiang, Jiangmeng Li, Lingyu Si, and Changwen Zheng. 2024. Rethinking Meta-Learning from a Learning Lens. *CoRR* abs/2409.08474 (2024). <https://doi.org/10.48550/ARXIV.2409.08474> arXiv:2409.08474
- [44] Jianhao Wang, Zhizhou Ren, Terry Liu, Yang Yu, and Chongjie Zhang. 2021. QPLEX: Duplex Dueling Multi-Agent Q-Learning. In *ICLR 2021: The Ninth International Conference on Learning Representations*.
- [45] Tonghan Wang, Jianhao Wang, Chongyi Zheng, and Chongjie Zhang. 2020. Learning Nearly Decomposable Value Functions Via Communication Minimization. In *ICLR 2020 : Eighth International Conference on Learning Representations*.
- [46] Bo Wu, Xiaoya Yang, Chuxiong Sun, Rui Wang, Xiaohui Hu, and Yan Hu. 2020. Learning effective value function factorization via attentional communication. In *2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 629–634.
- [47] Tete Xiao, Xiaolong Wang, Alexei A. Efros, and Trevor Darrell. 2021. What Should Not Be Contrastive in Contrastive Learning. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net. <https://openreview.net/forum?id=CZ8Y3NzuVzO>
- [48] Di Xue, Lei Yuan, Zongzhang Zhang, and Yang Yu. 2022. Efficient Multi-Agent Communication via Shapley Message Value. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, Lud De Raedt (Ed.). International Joint Conferences on Artificial Intelligence Organization, 578–584. <https://doi.org/10.24963/ijcai.2022/82> Main Track.
- [49] Denis Yarats, Ilya Kostrikov, and Rob Fergus. 2021. Image Augmentation Is All You Need: Regularizing Deep Reinforcement Learning from Pixels. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net. <https://openreview.net/forum?id=GY6-6sTvGaf>
- [50] Chao Yu, Akash Velu, Eugene Vinitisky, Jiaxuan Gao, Yu Wang, Alexandre Bayen, and Yi Wu. 2022. The surprising effectiveness of ppo in cooperative multi-agent games. *Advances in Neural Information Processing Systems* 35 (2022), 24611–24624.
- [51] Lei Yuan, Jianhao Wang, Fuxiang Zhang, Chenghe Wang, Zongzhang Zhang, Yang Yu, and Chongjie Zhang. 2022. Multi-agent incentive communication via decentralized teammate modeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 9466–9474.
- [52] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. 2021. Barlow Twins: Self-Supervised Learning via Redundancy Reduction. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 12310–12320. <http://proceedings.mlr.press/v139/zbontar21a.html>
- [53] Sai Qian Zhang, Qi Zhang, and Jieyu Lin. 2019. Efficient communication in multi-agent reinforcement learning via variance based control. In *Advances in Neural Information Processing Systems*. 3235–3244.
- [54] Sai Qian Zhang, Qi Zhang, and Jieyu Lin. 2020. Succinct and robust multi-agent communication with temporal message control. *Advances in Neural Information Processing Systems* 33 (2020), 17271–17282.