

Consensus Learning for Cooperative Multi-Agent Reinforcement Learning

Zhiwei Xu, Bin Zhang, Dapeng Li, Zeren Zhang, Guangchong Zhou, Hao Chen, Guoliang Fan

Institute of Automation, Chinese Academy of Sciences
School of Artificial Intelligence, University of Chinese Academy of Sciences
Beijing, China

{xuzhiwei2019, zhangbin2020, lidapeng2020, zhangzeren2021, zhoughuangchong2021, chenhao2019, guoliang.fan}@ia.ac.cn

Abstract

Almost all multi-agent reinforcement learning algorithms without communication follow the principle of centralized training with decentralized execution. During the centralized training, agents can be guided by the same signals, such as the global state. However, agents lack the shared signal and choose actions given local observations during execution. Inspired by viewpoint invariance and contrastive learning, we propose consensus learning for cooperative multi-agent reinforcement learning in this study. Although based on local observations, different agents can infer the same consensus in discrete spaces without communication. We feed the inferred one-hot consensus to the network of agents as an explicit input in a decentralized way, thereby fostering their cooperative spirit. With minor model modifications, our suggested framework can be extended to a variety of multi-agent reinforcement learning algorithms. Moreover, we carry out these variants on some fully cooperative tasks and get convincing results.

Introduction

Multi-agent reinforcement learning has received increasing attention recently due to its successful applications in games (Vinyals et al. 2019), transportation (Chu et al. 2020), and networks (Nasir and Guo 2019). And multi-agent cooperation has gradually gained popularity as a research field: how to control multiple agents to work together to solve the same problem simultaneously? Many multi-agent reinforcement learning algorithms based on centralized training with decentralized execution (CTDE) paradigms (Lowe et al. 2017) have been proposed, including multi-agent deep deterministic policy gradient (MADDPG) (Lowe et al. 2017), counterfactual multi-agent policy gradients (COMA) (Foerster et al. 2018), and numerous value decomposition methods (Sunehag et al. 2018; Rashid et al. 2018, 2020; Wang et al. 2020, 2021; Mahajan et al. 2019). All of the above methods use a centralized critic to address the instability problem in multi-agent systems.

Methods built on CTDE, however, have a flaw. While an agent can share the information with other agents during the centralized learning, it can only make decisions regarding its own individual observations for decentralized ex-



Figure 1: An example on the viewpoint invariance. In the StarCraft multi-agent challenge, agents receive local observations drawn within their field of view. Even though observations of all agents are completely different at each timestep, they nonetheless reflect the same global state.

ecution. The lack of common guidance is particularly severe in the partially observable Markov decision process (POMDP) (Monahan 1982). As a result, agents lack the desire to cooperate when executing. This issue can be resolved by using certain communication-based multi-agent reinforcement learning techniques (Sukhbaatar, Szlam, and Fergus 2016; Foerster et al. 2016). However, communication methods also bring about problems like choosing the information to pass and additional bandwidth requirements.

Inspired by computer vision (Morel and Yu 2009; Xia, Chen, and Aggarwal 2012), we propose **CO**nsensus **Le**arning (COLA) for multi-agent reinforcement learning based on the invariance to views. In POMDPs, despite the fact that local observations of each agent are unique, the observations of all agents at the same moment are different representations of the same global state, as shown in Figure 1. It is like a multi-view problem (Xu, Tao, and Xu 2013), where the same object can look different from various viewpoints. If the agents are aware of the invariance, they can explicitly select cooperative actions by inferring the same consensus on the state from their different local observations during the decentralized execution. Meanwhile, there is no communication throughout the entire process above.

There are several ways to represent the consensus of the agents on the state. The easiest way is to use the real state as consensus. However, it is not feasible in the implementation to accurately infer the global state from local observations. As a representation learning framework based on the similarity between samples, contrastive learning (Sermanet

et al. 2018; Wu et al. 2020; Khosla et al. 2020) has recently emerged as the most popular field in self-supervised learning. While ensuring that the learned representation can retain as much information as possible, contrastive learning promotes similar samples to be close together in the representation space. COLA utilizes contrastive learning to encode local observations into discrete latent spaces through the form of invariances. Even though the local observations between agents at the same moment may be quite dissimilar, we still treat them as similar when training the consensus representation model. Thus, even if agents in COLA receive entirely different local observations, they can still infer the same discrete state latent variables during execution. COLA enables agents to learn the “tacit agreement” in human social activities: although there is no substantial interaction between agents, they have reached a consensus on the current state.

We evaluate COLA on variants of the multi-agent particle environments (Lowe et al. 2017), the challenging micromanagement task of StarCraft II (Samvelyan et al. 2019), and the mini-scenarios of Google Research Football (Kurach et al. 2020). We demonstrate that COLA outperforms the previous baselines through experimental results. Furthermore, our approach of simply adding a one-hot consensus encoding to the network input can be extended to any other multi-agent reinforcement learning algorithm.

Related Work

Contrastive Learning

Many recent studies have focused on mining knowledge from data through self-supervised learning to improve the ability of models to solve downstream tasks. As a typical discriminative method, contrastive learning is easier to train and understand than other generative self-supervised learning methods (Kingma and Welling 2014; Goodfellow et al. 2014). A dataset in contrastive learning tasks is typically organized by similar and dissimilar pairs. Then contrastive learning methods are trained by reducing the distance between augmented views of the same sample and increasing it between different ones (Wang and Isola 2020). In this way, contrastive learning enables the representation model to learn the general knowledge contained in the dataset. However, the selection of dissimilar instances, or negative instances, can largely influence the performance of contrastive learning. This process often involves manual selection, limiting further development of contrastive learning. Some new work (Grill et al. 2020; Chen and He 2021; Zbon-tar et al. 2021) cancels the explicit introduction of negative examples by using techniques like clustering (Caron et al. 2020), but they can still get good performances and avoid collapse, a common failure in contrastive learning.

Contrastive Learning for Reinforcement Learning

Since contrastive learning can learn the underlying features of the data, it can be used to extract key information from the raw inputs in reinforcement learning tasks. By comparing states (often in pixel spaces) at different times, contrastive learning can map the original complex state to a low-dimensional latent space, removing redundant

information unrelated to learning. So this feature extraction step can improve the sample efficiency of single-agent reinforcement learning, especially in high-dimensional environments. Some studies (Laskin, Srinivas, and Abbeel 2020; Kostrikov, Yarats, and Fergus 2021; Laskin et al. 2020) successfully use contrastive learning as representation learning before reinforcement learning and achieve state-of-the-art data efficiency on pixel-based RL tasks. Another approach is Contrastive Predictive Coding (van den Oord, Li, and Vinyals 2018; Hénaff 2020), which uses contrastive learning as an auxiliary task to predict the future state conditioned on past observations and actions. In addition, contrastive learning can provide a reward function within a reinforcement learning system (Sermanet et al. 2018; Dwivedi et al. 2018).

Compared with single-agent reinforcement learning, there is less work in the multi-agent domain for contrastive learning. Liu, Yan, and Alahi (2021) induce the emergence of a common language by maximizing the mutual information between messages of a given trajectory in a contrastive learning manner. Lo and Sengupta (2022) propose a contrastive method for learning socially-aware motion representations and successfully applied it in human trajectory forecasting and crowd navigation algorithms. Inspired by invariance to views, our proposed COLA can lead agents to reach a tacit agreement in a multi-agent setup where agents only receive partial observations of the same state. To the best of our knowledge, this is the first study to introduce contrastive learning into the decentralized partially observable Markov decision process problem.

Preliminaries

Dec-POMDP and CTDE Methods

A fully cooperative multi-agent task can be viewed as a decentralized partially observable Markov decision process (Dec-POMDP) (Oliehoek and Amato 2016), which can be modeled as a tuple $G = \langle S, U, A, P, r, Z, O, n, \gamma \rangle$. $s \in S$ is the global state of the environment. In the CTDE principle, the global state s is accessible during the centralized training but not during the decentralized execution. At each time step, each agent $a \in A := \{1, \dots, n\}$ determines the appropriate action $u^a \in U$ conditioned on local observation $z^a \in Z$, which is provided by the observation function $O(s, a) : S \times A \rightarrow Z$. $\mathbf{u} \in \mathbf{U} \equiv U^n$ denotes the joint action of all agents. The state transition function is written as $P(s' | s, \mathbf{u}) : S \times \mathbf{U} \times S \rightarrow [0, 1]$. All agents in Dec-POMDPs share one reward function: $r(s, \mathbf{u}) : S \times \mathbf{U} \rightarrow \mathbb{R}$. The last term γ is the discount factor.

Then we introduce some CTDE methods to solve Dec-POMDP problems. A mixing network and several agent networks make up the framework of value decomposition methods. For agent a , we can get the individual action-value estimation $Q_a(\tau^a, \cdot)$ by feeding the local observations to its own value net. The mixing network decomposes the joint action-value function into values related to per-agent utilities. To guarantee the consistency between the local greedy actions and global ones, we make the following assumption known as the Individual-Global-Max (IGM) (Son et al. 2019):

$$\arg \max_{u^a} Q_{tot}(\boldsymbol{\tau}, \mathbf{u}) = \arg \max_{u^a} Q_a(\tau^a, u^a), \quad \forall a \in A,$$

where τ represents the joint action-observation histories. Value decomposition methods are trained end-to-end by minimizing the loss function:

$$\mathcal{L} = (y_{tot} - Q_{tot}(\tau, \mathbf{u}))^2, \quad (1)$$

where $y_{tot} = r + \gamma \max_{\mathbf{u}'} \hat{Q}_{tot}(\tau', \mathbf{u}')$. $\hat{Q}_{tot}(\cdot)$ is the target network of the joint action-value function. In the value decomposition method, the agent greedily chooses actions according to its individual action-value function Q_a during the decentralized execution.

MADDPG (Lowe et al. 2017) is an actor-critic method based on the CTDE paradigm for continuous action tasks. Each agent learns a deterministic policy μ_a in an off-policy manner. MADDPG consists of a centralized critic and several independent actors corresponding to each agent. During the centralized training phase, the critic can use global information to estimate the joint action-value $Q_a(s, u^1, \dots, u^n)$ for each agent a . We can train the critic by minimizing the TD-error like Eq. (1). The gradient of the individual deterministic policy μ_a can be derived as follow:

$$\nabla J(\mu_a) = \mathbb{E}_{\mathcal{D}} \left[\nabla \mu_a(\tau^a) \nabla_{u^a} Q_a(s, u^1, \dots, u^n) \Big|_{u^a = \mu_a(\tau^a)} \right], \quad (2)$$

where u^a in $Q_a(s, u^1, \dots, u^n)$ is sampled from the current policy μ_a , while the actions of other agents are sampled from the replay buffer \mathcal{D} .

Value decomposition methods and policy-based multi-agent methods are very different in structure. However, they both follow CTDE principles, and suffer from a lack of common guidance during the decentralized execution in fully cooperative tasks.

Knowledge Distillation with No Labels

Many early contrastive learning studies require a large batch pool or suitable negative samples, so they cannot be directly applied to reinforcement learning. In Knowledge Distillation with No Labels (DINO) (Caron et al. 2021) which is inspired by knowledge distillation (Hinton, Vinyals, and Dean 2015), a student network g_{θ_S} directly predicts the output of a teacher network g_{θ_T} to simplify self-supervised training. For two different augmentations of the same image, x and x' , two networks output probability distributions over K dimensions, represented as $P_T(x)$ and $P_S(x')$ respectively. $P_T(x)$ is given as follows:

$$P_T(x)^{(i)} = \frac{\exp(g_{\theta_T}(x)^{(i)}/\tau_T)}{\sum_{k=1}^K \exp(g_{\theta_T}(x)^{(k)}/\tau_T)}, \forall i \in \{1, \dots, K\}. \quad (3)$$

τ_T is the temperature coefficient that controls the sharpness of the output dimension. A similar formula holds for $P_S(x')$ with temperature τ_S . So in the case where no labels are available, the teacher network can give pseudo-labels to unlabelled data. The student network needs to match these pseudo-labels by minimizing the cross-entropy loss:

$$\min_{\theta_S} \sum_{\substack{x, x' \in \mathcal{V} \\ x \neq x'}} H(P_T(x), P_S(x')),$$

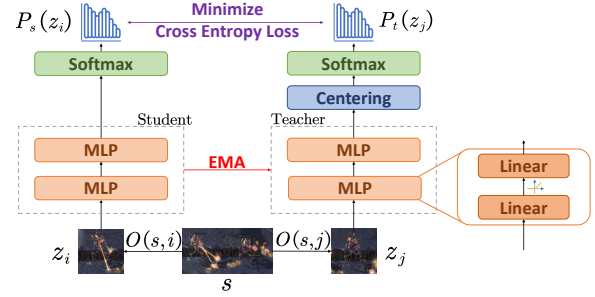


Figure 2: The overall architecture of the consensus builder.

where \mathcal{V} represents different views of an image and $H(a, b) = -a \log b$. Unlike general knowledge distillation methods, the structure of the teacher network in DINO is consistent with that of the student network, and its weight parameters also use an exponential moving average (EMA) on the student weights. Furthermore, to avoid trivial solutions, the centering operation prevents one dimension from dominating but encourages the output of the teacher network to the uniform distribution. So we obtain the modified output of the teacher network according to Eq. (3):

$$P_T(x) = \text{Softmax}((g_{\theta_T}(x) - \lambda)/\tau_T), \quad (4)$$

where the center λ is updated with an exponential moving average of $g_{\theta_T}(x)$. DINO applies the above framework to the Vision Transformer and obtains features that outperform those trained with labels.

In multi-agent systems, observations of each agent can be seen as different views of the global state. Agents in COLA can learn the identical discrete representation from these different local observations by employing contrastive learning. We believe the inferred consensus can ignite collaboration between agents during the decentralized execution.

Consensus Learning

In this section, we propose COLA, a novel multi-agent cooperative reinforcement learning framework that can explicitly guide agents to make cooperative decisions in decentralized execution. It is worth noting that our proposed method is universal and any CTDE multi-agent algorithm already in existence can be combined with COLA.

Consensus Builder

During the centralized training phase of the conventional CTDE methodology, each agent can improve coordination based on consistent environmental signals, such as the same global state. Even in decentralized execution, agents also need to derive shared signals based on their local information to get consistent guidance and inspiration. In a multi-agent system, while local observations are different for each agent, they all correspond to the same global state. This is similar to the multi-view problem in computer vision. Data augmentation (Shorten and Khoshgoftaar 2019) is a method that enables limited data to generate more equivalent data, hence increasing the diversity of training

data sets. We can treat observation function $O(s, a) \in Z$ of each agent as an augmentation operation. The global state s is data-augmented to form local observations z of each agent. However, it is not feasible to directly map local observations to a continuous low-dimensional space to obtain shared signals using general contrastive learning methods, because there are always subtle differences between the continuous signals inferred by different agents. Furthermore, since both consensus learning and reinforcement learning are carried out simultaneously in COLA, the best way to improve robustness and make learning easier is to map the consensus to a discrete space. To achieve this goal, we adopt a DINO-like framework to learn discrete consensus, named consensus builder. The goal of the consensus builder is to group local observations for the same state into the same class c . Then the class c is the discrete global consensus shared by all agents in that state.

Enlightened by DINO, we enable the consensus builder to simulate knowledge distillation through student and teacher networks as shown in Figure 2. The centering operation on the output of the teacher network avoids collapse. In multi-agent reinforcement learning, we naturally treat the observation function as data augmentation operations. In addition, since most of the observations in the multi-agent environments we used in this paper are low-dimensional and not complex high-dimensional signals like images, simple multi-layer perceptrons (MLP) are enough to be used to build the main body of the network. For the state s , $z^a = O(s, a) \in Z$ represents the local observation of the agent a . According to Eq. (3) and Eq. (4), we can obtain the probabilities $P_T(z^a)$ and $P_S(z^a)$ of the agent’s discrete consensus under the condition of local observation z , where $a \in \{1, 2, \dots, n\}$. We perform the pairwise comparison between the inferred consensus of all agents and minimize the cross-entropy loss:

$$\mathcal{L}_{CB} = \sum_{a,b} H(P_T(z^a), P_S(z^b)), \quad (5)$$

where $a, b \in \{1, 2, \dots, n\}$ and $a \neq b$. By minimizing \mathcal{L}_{CB} , consensus builders can provide identical categorical distributions for different views of the same state. $c^a \in \{1, 2, \dots, K\}$ represents the discrete consensus inferred by the agent a based on the observation z^a , which can be computed according to the output of the student network:

$$c^a = \arg \max_c P_S(z^a)^{(c)}. \quad (6)$$

Since Eq. (5) requires local observations of all agents and Eq. (6) only requires that of a single agent, the consensus builder also follows the paradigm of centralized training with decentralized execution. The center λ in Eq. (4) and the teacher network are updated with an exponential moving average. All hyperparameter settings and the details of the implementation can be found in Appendix B.

COLA Architecture

Because the consensus builder can judge the current consensus in decentralized execution, we can use it as additional




Figure 3: Illustration of multi-agent reinforcement learning methods with consensus learning. Left: The value decomposition method with consensus learning. Right: The multi-agent actor-critic method with consensus learning.

information for each agent’s decision-making. We concatenate local observations with the one-hot consensus and feed them into the network. This simple method can be easily deployed to various multi-agent reinforcement learning algorithms. We input the merged information into the agent network so that value decomposition methods can be trained end-to-end to minimize the following loss:

$$\mathcal{L}_{RL} = (y_{tot} - Q_{tot}(\tau, c, \mathbf{u}))^2, \quad (7)$$

where $y_{tot} = r + \gamma \max_{\mathbf{u}'} \hat{Q}_{tot}(\tau', c', \mathbf{u}')$. While for policy-based multi-agent reinforcement learning methods such as MADDPG, $Q_a(s, c^a, u^1, \dots, u^n)$ denotes the centralized action-value function of agent a . And the following policy gradient can be calculated:

$$\nabla J(\mu_a) = \mathbb{E}_{\mathcal{D}} \left[\nabla \mu_a(\tau^a, c^a) \nabla_{u^a} Q_a(s, c^a, u^1, \dots, u^n) \Big|_{u^a = \mu_a(\tau^a, c^a)} \right], \quad (8)$$

Figure 3 depicts the overview of the aforementioned multi-agent reinforcement learning methods with consensus learning. Since the inferred consensus is discrete, we do not propagate gradients about reinforcement learning through the consensus builder. As a result, the complexity of the reinforcement learning model was not significantly increased. So if COLA can improve the performance, we can clearly know that the improvement is brought by the consensus learning mechanism rather than by increasing model capacity. To prevent the sparse input due to the large number of consensus classes K , we add an embedding layer to map the originally one-hot hard consensus signal into a dense low-dimensional space. Although different local observations guide the agent to make different decisions, the same consensus given by the consensus builder based on invariance drives the collaboration of the agents. We argue that consensus learning can solve the problem of agents lacking cooperative signals when they execute in a decentralized way.

Experiment

We evaluate COLA in three challenging environments: the multi-agent particle environments (MPE), the StarCraft




Figure 4: Mean episode return on different tasks in the multi-agent particle environment.

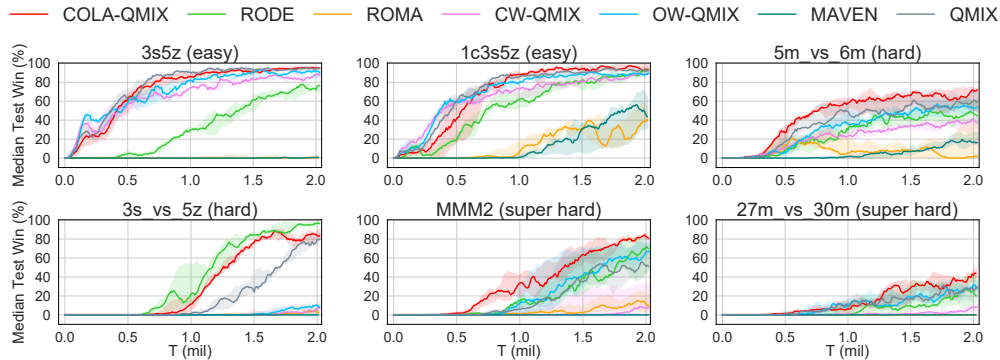


Figure 5: Performance comparison with baselines in different SMAC scenarios.

multi-agent challenge (SMAC), and Google Research Football (GRF). The detailed descriptions for the three environments can be found in Appendix A. We apply COLA to different CTDE algorithms for different environments and then compare them with baselines. Finally, we visualize the consensus inferred by agents. In this section, we present experimental results and answer the following three questions: (a) Can COLA be applied to various CTDE algorithms and improve their performance? (b) Is the superior performance of COLA due to consensus learning? (c) How does the hyperparameter K affect the performance of COLA?

Performance on Multi-Agent Particle Environments

There are different continuous action tasks in the multi-agent particle environments. We selected three representative cooperation scenarios: *Cooperative Predator-Prey*, *Cooperative Navigation*, and a new scenario we designed, *Cooperative Pantomime*. In *Cooperative Predator-Prey*, a single prey is the built-in heuristic agent, and we control a team of three predators. *Cooperative Pantomime* is similar to *Cooperative Communication* without the communication mechanism. The two agents in *Cooperative Pantomime* want to get to their target landmark, which is only known by another agent. So the agent can only infer the location of its goal from actions of another agent. Details of the environment can be found in Appendix A.1.

Since MADDPG often fails in complex environments like SMAC, we compare MADDPG with COLA-MADDPG in multi-agent particle environments, where COLA-MADDPG is obtained by combining the consensus builder with MADDPG. We carry out each experiment with 5 random seeds,

and the mean episode returns are shown with a 95% confidence interval. Since the scenarios in MPE is simple, we set the number of classes of consensus $K = 4$.

Figure 4 shows the learning curves of COLA-MADDPG and MADDPG in different scenarios. COLA-MADDPG can successfully solve the task in any scenario and performs better than MADDPG. Each agent in MADDPG corresponds to a policy network, which may not be conducive to encouraging agents to cooperate even in the same state. In COLA-MADDPG, we feed the signals inferred by the consensus builder as explicit input to each policy network, which helps the agents to have an abstract understanding of the state.

Performance on StarCraft II and Google Research Football

Many value decomposition methods have achieved impressive performance in complex scenarios such as StarCraft II and Google Research Football. We develop COLA-QMIX by combining COLA with the vanilla value decomposition method QMIX (Rashid et al. 2018). To demonstrate the superiority of our proposed COLA framework, we evaluate COLA-QMIX and baselines including QMIX, MAVEN (Mahajan et al. 2019), ROMA (Wang et al. 2020), Weighted QMIX (Rashid et al. 2020), and RODE (Wang et al. 2021) on SMAC and Google Research Football. To eliminate unfair comparisons like those present in the original literature of some baselines, we evaluate the algorithms in the exact same environmental settings and show their real performance.

SMAC is a novel multi-agent testbed containing various micro-management tasks and we selected several representative scenarios. The version of StarCraft we used is




Figure 6: Comparison of our approach against baselines on Google Research Football.




Figure 7: The bars represent the relative size of neural network models related to reinforcement learning for some value decomposition variants over QMIX.

SC2.4.6.2.69232 instead of the simpler SC2.4.10. The multi-agent cooperation problems in hard and super hard scenarios have a large state space, so we set $K = 32$. For the easy scenarios, however, we still keep $K = 4$. Furthermore, the number of surviving agents always changes in one episode. Local observations of dead agents are padded with zeros, which violates the viewpoint invariance principle. Therefore, we disregard local observations of dead agents when training the consensus builder.

The median win ratios are shown in Figure 5, and the 25-75% percentiles are shaded. In the homogeneous agent scenarios $5m_vs_6m$ and $27m_vs_30m$, COLA-QMIX improves the performance of vanilla QMIX and outperforms other methods. The *MMM2* scenario contains a wide variety of agents. COLA-QMIX can map the observations of different types of agents into a low-dimensional space so that the agents get the same cooperative signal. Besides, COLA still achieves slight performance improvement on easy scenarios, which is difficult for complex value decomposition variants. We also compare the performance of COLA-VDN and VDN in SMAC without access to the global state. The results in Appendix C show that the consensus inferred by COLA can indeed be used as a global shared signal like the state.

In Google Research Football, agents can be trained to play football. There are many mini-scenarios in Football Academy. To speed up the training, we take the goal of our players or the ball into our half as the basis for determining the termination of an episode. In addition to scoring rewards, the agents can receive rewards when the ball approaches the goal. Furthermore, there are many types of agents, including strikers and midfielders. We performed COLA-QMIX and other baselines in three scenarios on Google Research Football.

Figure 6 shows the learning curves of different algorithms in Google Research Football. We focus on comparing the performance of vanilla QMIX and its variant with consensus learning. *Academy_corner* is a challenging scenario where players require more coordination to score goals. The other players need to agree with the server in the corner to reduce the difficulty of scoring. COLA-QMIX outperforms other algorithms in almost all scenarios, especially in *academy_corner*.

In summary, we apply COLA to the multi-agent actor-critic algorithm and the value decomposition method, respectively. The resulting algorithms achieve convincing results on three different experimental platforms.

Case Study and Visualization

To explore what makes the COLA framework superior to the original algorithm, we first rule out the factor of increasing the capacity of the neural network model. Figure 7 shows the relative size of the reinforcement learning model for various value decomposition variants over QMIX. Since the gradient of reinforcement learning tasks in the COLA framework cannot be transferred to the consensus builder, the actual model size does not increase significantly. Some complex value decomposition variants significantly increase the capacity of the model. The experimental results, however, show that these variants do not result in an improvement proportionate to the size of their models. The COLA framework brings the most notable performance gains with the smallest model.

Then we provide a case study to explore the role of the consensus builder and visualize its impact in an intuitive form. We perform our research on the *MMM2* scenario from the SMAC benchmark. As shown in Figure 8, in the 2D t-SNE (van der Maaten and Hinton 2008) embedding of states, the states with the same consensus generated by the agent tend to cluster together. Besides, states with the same consensus often have similar estimated values or true returns. It should be noted that the consensus builder learns the discrete consensus only based on the given observations, and the state information is not involved in training. So we can conclude that although the consensus output by the consensus builder is directly determined by the observations of each agent, it can reflect the global state. In other words, the consensus inferred by each agent during the decentralized execution can be used as a shared cooperation signal similar to the global state. Furthermore, when sharing different consensus, the agents implement different global strategies. By analyzing and visualizing

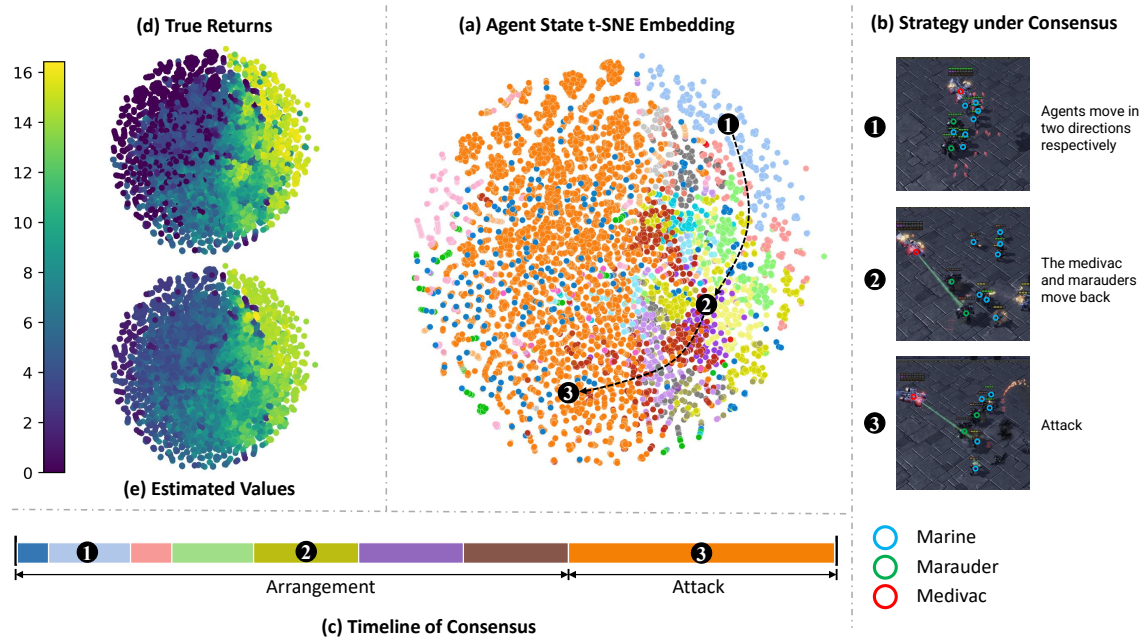


Figure 8: The analysis and visualization of agents’ consensus in *MMM2* scenario. (a) The 2D t-SNE embedding of states. Each point is colored according to the consensus inferred by the agents based on local observations in this state. (b) The global strategies of the agents under the guidance of different consensus. (c) The change of consensus in an episode. The consensus represented by each color is consistent with that in (a). (d) The true returns of agents and (e) the estimated values.

consensuses, we can point out that the consensus signal inferred by the consensus builder is beneficial for agents to cooperate explicitly during the decentralized execution.

Ablation Study

In this section, we will investigate how the amount of consensus K affects the performance of COLA. We conduct an ablation study on two scenarios with various difficulties in SMAC. While keeping other settings unchanged, we compare the performance of COLA-QMIX under different values of K . Figure 9 displays the different results in the easy scenario *3s5z* and the super hard scenario *MMM2*. In the *3s5z* scenario, COLA-QMIX performs the worst when $K = 32$ but learns faster when K takes relatively tiny values. In the *MMM2* scenario, conversely, COLA-QMIX performs significantly better when $K = 32$ or $K = 16$ than when K is small. Therefore, we conclude that the choice of K depends on the difficulty of the task. We should pick a higher K value when the reinforcement learning task is complex and vice versa.

Conclusion

Inspired by viewpoint invariance, We propose a consensus learning framework called COLA. COLA can infer a shared consensus based on local observations of agents to alleviate the lack of the same guidance during the decentralized execution. We believe that the same consensus can encourage agents to act cooperatively. To the best of our knowledge, COLA is the first work to address the Dec-POMDP problem by integrating contrastive learning into




Figure 9: Influence of the K for COLA-QMIX.

multi-agent reinforcement learning. We combine consensus learning with various multi-agent algorithms and evaluate their performances in several environments. Experimental results demonstrate that our proposed COLA can improve the performance of the original algorithms on fully cooperative tasks. We believe the COLA framework is the most cost-effective, bringing remarkable performance improvement with minor changes of reinforcement learning models. Further research on how to choose the hyperparameter K would be of interest.

Acknowledgments

The work is supported by the National Defence Foundation Reinforcement Fund.

References

- Caron, M.; Misra, I.; Mairal, J.; Goyal, P.; Bojanowski, P.; and Joulin, A. 2020. Unsupervised Learning of Visual Features by Contrasting Cluster Assignments. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging Properties in Self-Supervised Vision Transformers. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 9630–9640.
- Chen, X.; and He, K. 2021. Exploring Simple Siamese Representation Learning. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 15745–15753.
- Chu, T.; Wang, J.; Codecà, L.; and Li, Z. 2020. Multi-Agent Deep Reinforcement Learning for Large-Scale Traffic Signal Control. *IEEE Transactions on Intelligent Transportation Systems*, 21: 1086–1095.
- Dwibedi, D.; Tompson, J.; Lynch, C.; and Sermanet, P. 2018. Learning Actionable Representations from Visual Observations. *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1577–1584.
- Foerster, J. N.; Assael, Y. M.; de Freitas, N.; and Whiteson, S. 2016. Learning to Communicate with Deep Multi-Agent Reinforcement Learning. In Lee, D. D.; Sugiyama, M.; von Luxburg, U.; Guyon, I.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, 2137–2145.
- Foerster, J. N.; Farquhar, G.; Afouras, T.; Nardelli, N.; and Whiteson, S. 2018. Counterfactual Multi-Agent Policy Gradients. In McIlraith, S. A.; and Weinberger, K. Q., eds., *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18), the 30th Innovative Applications of Artificial Intelligence Conference (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, 2974–2982. AAAI Press.
- Goodfellow, I. J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A. C.; and Bengio, Y. 2014. Generative Adversarial Nets. In Ghahramani, Z.; Welling, M.; Cortes, C.; Lawrence, N. D.; and Weinberger, K. Q., eds., *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, 2672–2680.
- Grill, J.; Strub, F.; Altché, F.; Tallec, C.; Richemond, P. H.; Buchatskaya, E.; Doersch, C.; Pires, B. Á.; Guo, Z.; Azar, M. G.; Piot, B.; Kavukcuoglu, K.; Munos, R.; and Valko, M. 2020. Bootstrap Your Own Latent - A New Approach to Self-Supervised Learning. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Hénaff, O. J. 2020. Data-Efficient Image Recognition with Contrastive Predictive Coding. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, 4182–4192. PMLR.
- Hinton, G. E.; Vinyals, O.; and Dean, J. 2015. Distilling the Knowledge in a Neural Network. *ArXiv*, abs/1503.02531.
- Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised Contrastive Learning. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Kingma, D. P.; and Welling, M. 2014. Auto-Encoding Variational Bayes. In Bengio, Y.; and LeCun, Y., eds., *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*.
- Kostrikov, I.; Yarats, D.; and Fergus, R. 2021. Image Augmentation Is All You Need: Regularizing Deep Reinforcement Learning from Pixels. *ArXiv*, abs/2004.13649.
- Kurach, K.; Raichuk, A.; Stanczyk, P.; Zajac, M.; Bachem, O.; Espeholt, L.; Riquelme, C.; Vincent, D.; Michalski, M.; Bousquet, O.; and Gelly, S. 2020. Google Research Football: A Novel Reinforcement Learning Environment. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, 4501–4510. AAAI Press.
- Laskin, M.; Lee, K.; Stooke, A.; Pinto, L.; Abbeel, P.; and Srinivas, A. 2020. Reinforcement Learning with Augmented Data. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Laskin, M.; Srinivas, A.; and Abbeel, P. 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*, volume 119 of *Proceedings of Machine Learning Research*, 5639–5650. PMLR.
- Liu, Y.; Yan, Q.; and Alahi, A. 2021. Social NCE: Contrastive Learning of Socially-aware Motion Representations. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 15098–15109.
- Lo, Y. L.; and Sengupta, B. 2022. Learning to Ground Decentralized Multi-Agent Communication with Contrastive Learning. *ArXiv*, abs/2203.03344.
- Lowe, R.; Wu, Y.; Tamar, A.; Harb, J.; Abbeel, P.; and Mordatch, I. 2017. Multi-Agent Actor-Critic for Mixed

- Cooperative-Competitive Environments. In Guyon, I.; von Luxburg, U.; Bengio, S.; Wallach, H. M.; Fergus, R.; Vishwanathan, S. V. N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, 6379–6390.
- Mahajan, A.; Rashid, T.; Samvelyan, M.; and Whiteson, S. 2019. MAVEN: Multi-Agent Variational Exploration. In Wallach, H. M.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E. B.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, 7611–7622.
- Monahan, G. E. 1982. State of the Art—A Survey of Partially Observable Markov Decision Processes: Theory, Models, and Algorithms. *Management Science*, 28: 1–16.
- Morel, J.-M.; and Yu, G. 2009. ASIFT: A New Framework for Fully Affine Invariant Image Comparison. *SIAM J. Imaging Sci.*, 2: 438–469.
- Nasir, Y. S.; and Guo, D. 2019. Multi-Agent Deep Reinforcement Learning for Dynamic Power Allocation in Wireless Networks. *IEEE Journal on Selected Areas in Communications*, 37: 2239–2250.
- Oliehoek, F.; and Amato, C. 2016. A Concise Introduction to Decentralized POMDPs. In *SpringerBriefs in Intelligent Systems*.
- Rashid, T.; Farquhar, G.; Peng, B.; and Whiteson, S. 2020. Weighted QMIX: Expanding Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. *arXiv: Learning*.
- Rashid, T.; Samvelyan, M.; de Witt, C. S.; Farquhar, G.; Foerster, J. N.; and Whiteson, S. 2018. QMIX: Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. In Dy, J. G.; and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, 4292–4301. PMLR.
- Samvelyan, M.; Rashid, T.; Witt, C. S. D.; Farquhar, G.; Nardelli, N.; Rudner, T. G. J.; Hung, C.-M.; Torr, P. H. S.; Foerster, J. N.; and Whiteson, S. 2019. The StarCraft Multi-Agent Challenge. *ArXiv*, abs/1902.04043.
- Sermanet, P.; Lynch, C.; Chebotar, Y.; Hsu, J.; Jang, E.; Schaal, S.; and Levine, S. 2018. Time-Contrastive Networks: Self-Supervised Learning from Video. *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 1134–1141.
- Shorten, C.; and Khoshgoftaar, T. M. 2019. A survey on Image Data Augmentation for Deep Learning. *Journal of Big Data*, 6: 1–48.
- Son, K.; Kim, D.; Kang, W. J.; Hostallero, D. E.; and Yi, Y. 2019. Qtran: Learning to factorize with transformation for cooperative multi-agent reinforcement learning. In *International conference on machine learning*, 5887–5896. PMLR.
- Sukhbaatar, S.; Szlam, A.; and Fergus, R. 2016. Learning Multiagent Communication with Backpropagation. In Lee, D. D.; Sugiyama, M.; von Luxburg, U.; Guyon, I.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, 2244–2252.
- Sunehag, P.; Lever, G.; Gruslys, A.; Czarnecki, W. M.; Zambaldi, V. F.; Jaderberg, M.; Lanctot, M.; Sonnerat, N.; Leibo, J. Z.; Tuyls, K.; and Graepel, T. 2018. Value-Decomposition Networks For Cooperative Multi-Agent Learning. *ArXiv*, abs/1706.05296.
- van den Oord, A.; Li, Y.; and Vinyals, O. 2018. Representation Learning with Contrastive Predictive Coding. *ArXiv*, abs/1807.03748.
- van der Maaten, L.; and Hinton, G. E. 2008. Visualizing Data using t-SNE. *Journal of Machine Learning Research*, 9: 2579–2605.
- Vinyals, O.; Babuschkin, I.; Czarnecki, W. M.; Mathieu, M.; Dudzik, A.; Chung, J.; Choi, D. H.; Powell, R.; Ewalds, T.; Georgiev, P.; Oh, J.; Horgan, D.; Kroiss, M.; Danihelka, I.; Huang, A.; Sifre, L.; Cai, T.; Agapiou, J. P.; Jaderberg, M.; Vezhnevets, A. S.; Leblond, R.; Pohlen, T.; Dalibard, V.; Budden, D.; Sulsky, Y.; Molloy, J.; Paine, T. L.; Gulcehre, C.; Wang, Z.; Pfaff, T.; Wu, Y.; Ring, R.; Yogatama, D.; Wünsch, D.; McKinney, K.; Smith, O.; Schaul, T.; Lillicrap, T. P.; Kavukcuoglu, K.; Hassabis, D.; Apps, C.; and Silver, D. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 1–5.
- Wang, T.; Dong, H.; Lesser, V. R.; and Zhang, C. 2020. ROMA: Multi-Agent Reinforcement Learning with Emergent Roles. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, 9876–9886. PMLR.
- Wang, T.; Gupta, T.; Mahajan, A.; Peng, B.; Whiteson, S.; and Zhang, C. 2021. RODE: Learning Roles to Decompose Multi-Agent Tasks. *ArXiv*, abs/2010.01523.
- Wang, T.; and Isola, P. 2020. Understanding Contrastive Representation Learning through Alignment and Uniformity on the Hypersphere. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, 9929–9939. PMLR.
- Wu, M.; Zhuang, C.; Mosse, M.; Yamins, D. L. K.; and Goodman, N. D. 2020. On Mutual Information in Contrastive Learning for Visual Representations. *ArXiv*, abs/2005.13149.
- Xia, L.; Chen, C.-C.; and Aggarwal, J. K. 2012. View invariant human action recognition using histograms of 3D joints. *2012 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, 20–27.
- Xu, C.; Tao, D.; and Xu, C. 2013. A Survey on Multi-view Learning. *ArXiv*, abs/1304.5634.
- Zbontar, J.; Jing, L.; Misra, I.; LeCun, Y.; and Deny, S. 2021. Barlow Twins: Self-Supervised Learning via Redundancy Reduction. In *ICML*.