

SIDE: State Inference for Partially Observable Cooperative Multi-Agent Reinforcement Learning

Zhiwei Xu, Yunpeng Bai, Dapeng Li, Bin Zhang, Guoliang Fan
 Institute of Automation, Chinese Academy of Sciences
 School of Artificial Intelligence, University of Chinese Academy of Sciences
 Beijing, China

{xuzhiwei2019,baiyunpeng2020,lidapeng2020,zhangbin2020,guoliang.fan}@ia.ac.cn

ABSTRACT

As one of the solutions to the decentralized partially observable Markov decision process (Dec-POMDP) problems, the value decomposition method has achieved significant results recently. However, most value decomposition methods require the fully observable state of the environment during training, but this is not feasible in some scenarios where only incomplete and noisy observations can be obtained. Therefore, we propose a novel value decomposition framework, named State Inference for value DEcomposition (SIDE), which eliminates the need to know the global state by simultaneously seeking solutions to the two problems of optimal control and state inference. SIDE can be extended to any value decomposition method to tackle partially observable problems. By comparing with the performance of different algorithms in StarCraft II micromanagement tasks, we verified that though without accessible states, SIDE can infer the current state that contributes to the reinforcement learning process based on past local observations and even achieve superior results to many baselines in some complex scenarios.

KEYWORDS

Reinforcement Learning; Variational Inference; Multi-Agent Learning; Graph Neural Networks

ACM Reference Format:

Zhiwei Xu, Yunpeng Bai, Dapeng Li, Bin Zhang, Guoliang Fan. 2022. SIDE: State Inference for Partially Observable Cooperative Multi-Agent Reinforcement Learning. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Online, May 9–13, 2022, IFAA-MAS*, 9 pages.

1 INTRODUCTION

Deep reinforcement learning has recently made breakthroughs in complex scenarios such as Atari games [26], robot control [29], and autonomous driving [3]. In the real world, however, due to the complexity of entity attributes and local observability, it is often impossible to obtain an effective representation for the state of the environment, which has a catastrophic impact on reinforcement learning. Different from the single-agent reinforcement learning tasks, there are multiple entities in the multi-agent system (MAS), so whether the state representation is appropriate plays a more important role in multi-agent reinforcement learning problems.

One of the traditional methods is to simply stack the local observations of all agents as the current state representation, but the most direct drawback of this is that as the number of agents increases, the dimension of the state representation will also increase dramatically.

On the one hand, to alleviate the partially observable problems, a lot of work has been proposed. Some model-based reinforcement learning algorithms [2] use environmental dynamics information to solve the POMDP problems, but transition functions are typically not available in most real tasks. Besides, because the recurrent neural network has the characteristics of integrating historical information, it was introduced into the vanilla reinforcement learning algorithms [8, 40, 44]. In this way, the performance has been improved while the algorithm remains model-free, even if the training of the recurrent neural network may require more trajectories of experience. Besides, there are state estimation methods based on variational inference [7, 9, 10] or belief tracking methods based on particle filters [13, 14, 23], but these approaches may not be practical in multi-agent cases. On the other hand, some work [12, 19] promotes the neural network to extract helpful information from the state of the complex environment by adding auxiliary tasks mainly including predicting the state of the next timestep. Intuitively, the key limitation of these studies is that they do not take the cases that the current states are unobserved into account.

As a notorious problem in MAS, Dec-POMDP [28] describes some collaboration problems. Since the global reward function is shared, it is necessary to allocate credit to each agent. In recent years, research on value decomposition has become very popular because of its simple implementation and excellent performance. The earliest value decomposition method is Value Decomposition Network (VDN) [38] and is followed by QMIX [33], one of the most popular multi-agent algorithms. Furthermore, many excellent variants of value decomposition methods [25, 39, 41, 42] have been recently proposed. However, most value decomposition methods use the global state by default during centralized training, which is not allowed in some environments where the global state is inaccessible.

In this paper, we propose State Inference for value DEcomposition (SIDE), a state variational estimation framework based on multi-agent value decomposition reinforcement learning methods. SIDE does not require dynamics information of the environment. It uses the variational graph auto-encoder [17] to integrate the local observation of all agents and reconstructs the state while reducing the dimension of the state space. Reinforcement learning and state inference in SIDE are carried out simultaneously, so it can promote the construction of states that is beneficial to maximizing returns.

Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online. © 2022 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

And because SIDE can be seen as a state inference mechanism based on the value decomposition method, it can be applied to any QMIX-style algorithm. We evaluated the performance of SIDE in the StarCraft II Dec-POMDP micromanagement tasks. The experimental results prove that when the true global state is unknown, the performance of SIDE can still be close to or even better than the methods of using states which are set manually.

2 RELATED WORK

In this section, we will summarize some recent work on partial observable problems in the case of single-agent and multi-agent respectively.

Single-Agent In a POMDP the agent additionally has perceptual uncertainty. Agents usually use past local observations and actions to understand the current state. This understanding is called belief. One of the most popular methods is the belief tracking method [13, 14, 23] based on particle filters, but this method often requires the preset model and state representation. Some similar work is to integrate information at different timesteps by introducing recurrent neural networks [8, 44], thereby allowing us to loosen the assumptions of earlier model-based methods and apply a model-free technique instead. The other is the state generation models [7, 10] based on variational inference. The generative model can abstract observations to output an effective state in compact latent spaces by optimizing the evidence lower bound. Sequential variational soft Q-learning network (SVQN) [9] represents POMDPs as a probabilistic graphical model (PGM) [18] and combines it with variational inference. Unlike the standard variational auto-encoder (VAE) [16], the distribution of state hidden variables is conditional on previous hidden states, so SVQN is equipped with an additional generative model to learn the conditional prior of hidden states. However, SVQN is only applicable to single-agent environments and there are still some interesting problems to be addressed in multi-agent systems.

Multi-Agent In the multi-agent environment, in addition to the partial observability of the environment itself, because each agent sometimes cannot observe other agents, this makes the problem of environmental uncertainty more complicated. Communication [15, 32, 37] can build a channel of information between various agents, thereby alleviating partial observable problems. Moreover, the framework of centralized training with decentralized execution (CTDE) allows the agent to share all the information with other agents during training. Much excellent work [4, 11, 22] has been proposed based on this paradigm and one of the most representative studies is the value decomposition method. The main idea of QMIX, a popular value decomposition method, is to input the action-values of all agents into the mixing network and then output the global action-value. The parameters of the mixing network are the output of a few hypernetworks [5] whose input is the global state. Although the above methods have effects in the multi-agent partially observable environments, they can only alleviate the problem of incomplete information between agents, rather than the invisible information problem of the environment itself. There are some prior studies [27, 43] using generative models and variational inference to learn belief representations in multi-agent systems. However,

they are not applicable to the Dec-POMDP problems and often require the information from the adversary.

Enlightened by SVQN, we have perfected the reconstruction of unobserved information in the multi-agent environment itself, that is, using variational inference with past local observations. Besides, we also confirmed whether QMIX can still outperform VDN under the premise of obtaining the same amount of information as VDN. Finally, it is proved according to the experimental results that in different scenarios, compared with QMIX under different state representation definitions, our proposed SIDE can find the most appropriate state representation. To illustrate it, we carry out the visualization of the embedding of latent states learned with our proposed method. The visualization results intuitively show that the latent state variables learned by SIDE are conducive to reinforcement learning.

3 BACKGROUND

3.1 Dec-POMDP

The Dec-POMDP is defined as a tuple $\langle S, \mathcal{U}, \mathcal{P}, r, \mathcal{Z}, O, n, \gamma \rangle$. $s \in S$ represents the state of the environment. At each timestep, each agent $a \in \mathcal{A} := \{1, \dots, n\}$ will take an action $u_a \in \mathcal{U}$ and then the local observation of each agent $z_a \in \mathcal{Z}$ is obtained by $O(s, a) : S \times \mathcal{A} \rightarrow \mathcal{Z}$. The actions of all agents form the joint action $\mathbf{u} \in \mathcal{U} \equiv \mathcal{U}^n$. The state transition function, which generates the next state of the environment, is defined as $\mathcal{P}(s'|s, \mathbf{u}) : S \times \mathcal{U} \times S \rightarrow [0, 1]$. In Dec-POMDPs, a common joint reward function $r(s, \mathbf{u}) : S \times \mathcal{U} \rightarrow \mathbb{R}$ is provided for all agents. $\gamma \in [0, 1)$ is the discount factor. The goal is to maximise the discounted return $R^t = \sum_{l=0}^{\infty} (\gamma)^l r^{t+l}$ in Dec-POMDPs.

In one episode, each agent will get the respective action-observation history $\tau_a \in T \equiv (\mathcal{Z} \times \mathcal{U})$. $\pi_a(u_a|\tau_a) : T \times \mathcal{U} \rightarrow [0, 1]$ denotes the policy of each agent. Note that for the convenience of presentation in this paper, the superscript represents the time and the subscript represents the identity number of the agent.

3.2 Value Decomposition

In the field of cooperative multi-agent reinforcement learning, it is not feasible to train each agent individually or treat all agents as a whole for joint training when the number of agents is large enough. Therefore, some research has proposed various methods between the above two extreme methods such as VDN, QMIX. These value decomposition methods try to achieve automated learning of decomposition of the joint value function based on the Individual-Global-Max (IGM) [36], where IGM assumes that the optimality of each agent is consistent with the optimality of all agents. The equation that describes IGM is as follows:

$$\arg \max_{\mathbf{u}} Q_{\text{tot}}(\boldsymbol{\tau}, \mathbf{u}) = \begin{pmatrix} \arg \max_{u_1} Q_1(\tau_1, u_1) \\ \vdots \\ \arg \max_{u_n} Q_n(\tau_n, u_n) \end{pmatrix},$$

where $\boldsymbol{\tau} \in T^n$ represents the joint action-observation histories of all agents. Q_{tot} is global action-value function and Q_a is the individual ones.

VDN assumes that the joint value function is linearly decomposable. However, the linear assumption is too simple to fit most

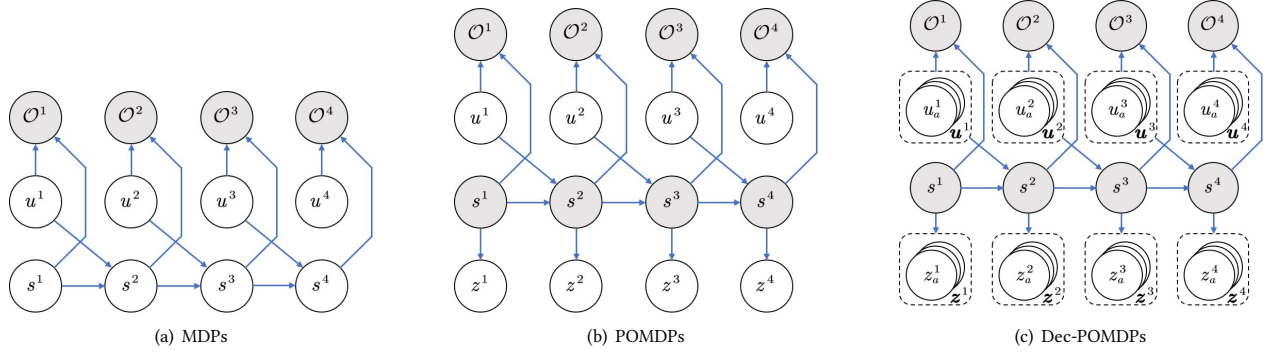


Figure 1: The graphical models for MDPs(a), POMDPs(b) and Dec-POMDPs(c). Grey nodes are hidden and white nodes are observable.

scenarios. Therefore, QMIX introduces a nonlinear global value function and assumes that the joint action-value function Q_{tot} is monotonic to the individual action-value function Q_a in order to satisfy the IGM assumption. In addition, many studies have presented interesting varieties of value decomposition and focused on different problems.

3.3 Probabilistic Graphical Models

The probabilistic graphical model (PGM) [20] is studied on the basis of probability theory and graph theory. It visualizes the probability model through the structure of the graph, allowing us to understand the relationship between variables in a complex distribution. In [21], the MDP problem is embedded in the PGM framework, so that the reinforcement learning problems can be viewed from another perspective. According to the paradigm of PGM, given the state $s \in \mathcal{S}$ and the action $u \in \mathcal{U}$, [21] introduced the optimality variable O to solve the optimality of MDP control problems. This variable O is a binary random variable. When $O^t = 1$, it means that the system at timestep t is optimal. On the contrary, when $O^t = 0$ means it is not optimal. Define the distribution of the variable O as:

$$p(O^t = 1 | s^t, u^t) \propto \exp(r(s^t, u^t)).$$

Then its variational evidence lower bound can be obtained as:

$$\log p(O^{1:T}) \geq \mathbb{E}_{(s^{1:T}, u^{1:T}) \sim \pi(s^{1:T}, u^{1:T})} \left[\sum_{t=1}^T r(s^t, u^t) - \log \pi(u^t | s^t) \right], \quad (1)$$

where $\pi(u^t | s^t)$ is the policy function. Therefore, increasing the optimal probability of the entire MDP $p(O^{1:T})$ is to optimize its optimal lower bound, which leads to the maximum entropy reinforcement learning [6]. Figure 1(a) shows the PGM of the ordinary MDP.

SVQN [9] combines latent state inference with maximum entropy reinforcement learning under a probabilistic graphical model, and optimizes these two modules simultaneously to address partial observation problems. Since the state s cannot be observed in the POMDP, it needs to be inferred from the local observation z and the action u . Such cases are depicted in Figure 1(b). SVQN

proposes a variational evidence lower bound that is different from the Equation (1):

$$\begin{aligned} & \log p(O^{0:T}, u^{0:T}, z^{1:T}) \\ &= \log \mathbb{E}_{q_\theta(s^{1:T} | O^{1:T}, u^{0:T}, z^{1:T})} \left[\frac{p(s^{1:T}, O^{0:T}, u^{0:T}, z^{1:T})}{q_\theta(s^{1:T} | O^{0:T}, u^{0:T}, z^{1:T})} \right] \\ &\geq \mathbb{E}_{q_\theta(s^{1:T} | O^{1:T}, u^{0:T}, z^{1:T})} \log \left[\frac{p(s^{1:T}, O^{0:T}, u^{0:T}, z^{1:T})}{q_\theta(s^{1:T} | O^{0:T}, u^{0:T}, z^{1:T})} \right] \\ &= \mathbb{E}_{q_\theta(s^{1:T} | O^{1:T}, u^{0:T}, z^{1:T})} \sum_{t=1}^T \left\{ r(s^t, u^t) + \log [p(u^t) p(z^t | s^t)] \right. \\ &\quad \left. - \mathcal{D}_{KL} [q_\theta(s^t | s^{t-1}, u^{t-1}, z^t) \| p(s^t | s^{t-1}, u^{t-1})] \right\}, \quad (2) \end{aligned}$$

where $p(u^t)$ is the action prior, for simplicity it is often regarded as a uniform distribution and it has nothing to do with the optimization of ELBO. And θ is the parameter of the approximate function $q(\cdot)$. $p(z^t | s^t)$ implies that the state s^t can generate the current observation z^t . For the last term KL divergence, since $p(s^t | s^{t-1}, u^{t-1})$ is unknown, SVQN presents an additional generative model $p_\theta^{\text{prior}}(s^t | s^{t-1}, u^{t-1})$ to approach the real distribution $p(s^t | s^{t-1}, u^{t-1})$. The structured variational inference is used to optimize the ELBO, that is, SVQN solves the two problems of optimal control and state inference by jointly optimizing the loss functions of two variational auto-encoders and the soft Q-learning.

3.4 Variational Graph Auto-Encoders

The variational graph auto-encoder (VGAE) [17], a framework that combines both variational auto-encoders and graph networks, is increasingly being used in graph structure data. Given a graph \mathcal{G} , its adjacency matrix is \mathbf{A} , and the node feature matrix is \mathbf{X} . The simplest VGAE is composed of a graph convolutional network (GCN) [1] encoder and an inner product decoder. First, each node i is mapped to a random variable distribution $q(z_i | \mathbf{X}, \mathbf{A})$ through GCN. After that, the latent variable representation z_i of each node is obtained by the reparameterization trick [16]. The decoder treats

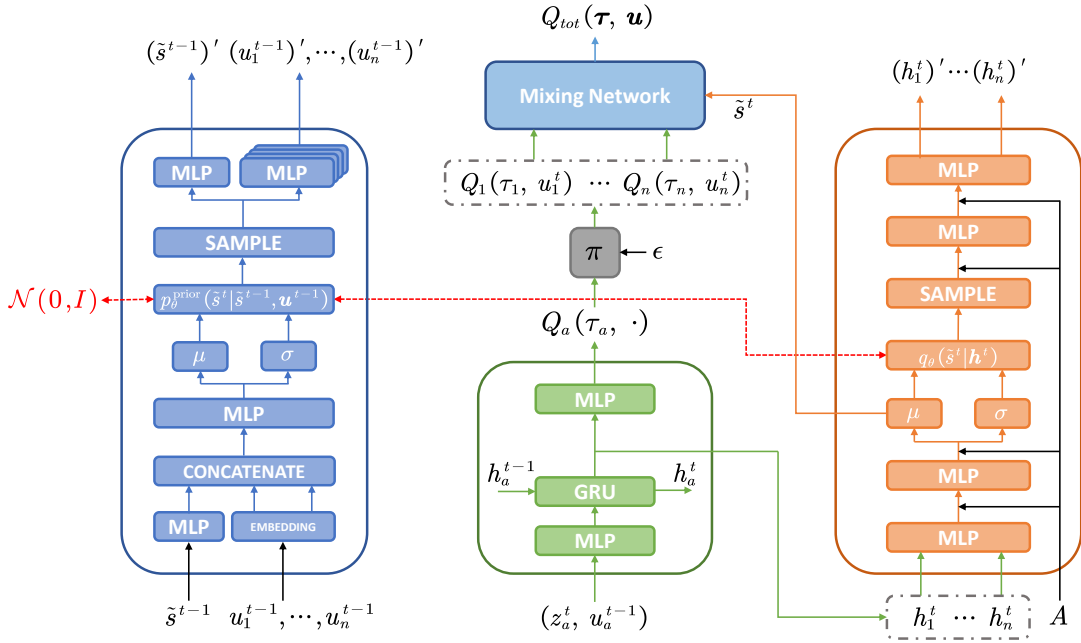


Figure 2: The overall architecture of SIDE. The left part is the prior function $p_\theta^{\text{prior}}(\cdot)$ and the state inference function $q_\theta(\cdot)$ based on the VGAE is on the right. The red dotted line connects the two distributions that need to calculate the KL divergence.

the inner product $A' = \sigma(ZZ^T)$ with respect to the hidden variables of all nodes as the reconstructed adjacency matrix, where $\sigma(\cdot)$ generally indicates the logistic sigmoid function. The optimization goal of VGAE is to minimize the difference between the real adjacency matrix and the reconstructed one, as well as the KL divergence between the random latent variable distribution and the prior distribution $p(\mathbf{Z})$ (usually set to a standard Gaussian distribution). The loss function can be written as:

$$\mathcal{L} = \mathbb{E}_{q(\mathbf{Z}|\mathbf{X},\mathbf{A})} [\log p(\mathbf{A}' | \mathbf{Z})] - \text{KL}[q(\mathbf{Z} | \mathbf{X}, \mathbf{A}) \| p(\mathbf{Z})].$$

Besides, many variations [30, 34] of the variational graph auto-encoder have been proposed recently. In order to reconstruct the node features \mathbf{X} instead of the adjacency matrix \mathbf{A} , the decoder in Graph convolutional Autoencoder using Laplacian smoothing and sharpening (GALA) [31] also uses a GCN structure that is symmetrical to the encoder. Contrary to the encoder based on Laplacian smoothing, the decoder is built on the basis of Laplacian sharpening. Compared with other variational graph auto-encoder variants, GALA shows superb results. In this paper, we will use GALA to integrate the local observations of all agents to reconstruct the state.

4 STATE INFERENCE FOR VALUE DECOMPOSITION

In this section, we will elaborate on the proposed algorithm State Inference for value DEcomposition (SIDE). At first we embed the Dec-POMDP into the PGMs and obtain the corresponding evidence lower bound, and then build a specific neural network framework.

4.1 Variational Lower Bound For Dec-POMDPS

Similar to the single-agent POMDP problem, in Dec-POMDP, the optimal variable \mathcal{O} and the global state s are both hidden variables that cannot be observed. However, the actions and observation variables of each agent in Dec-POMDP are different. Since the reward function is shared by all agents and \mathcal{O} is related to the reward r , all agents also share an optimal variable \mathcal{O} in PGMs. The PGMs framework of Dec-POMDP is shown in Figure 1(c).

First, we construct an inference function $q_\theta(\cdot)$ for the latent state, where θ is the learnable parameter. Since we need to solve the two problems of optimal control and state inference at the same time, here we use structured variational inference to optimize the variational evidence lower bound of Dec-POMDP, where the optimal policy we use the function $q_\pi(\cdot)$ to approach. It should be noted that, in order to facilitate comparison, we use the vanilla q-learning method instead of the soft q-learning method to train $q_\pi(\cdot)$. From above we can deduce the evidence lower bound of Dec-POMDP:

$$\begin{aligned} & \log p(\mathcal{O}^{0:T}, \mathbf{u}^{0:T}, \mathbf{z}^{1:T}) \\ &= \log \mathbb{E}_{q_\theta(s^{1:T} | \mathcal{O}^{1:T}, \mathbf{u}^{0:T}, \mathbf{z}^{1:T})} \left[\frac{p(s^{1:T}, \mathcal{O}^{0:T}, \mathbf{u}^{0:T}, \mathbf{z}^{1:T})}{q_\theta(s^{1:T} | \mathcal{O}^{0:T}, \mathbf{u}^{0:T}, \mathbf{z}^{1:T})} \right] \\ &\geq \mathbb{E}_{q_\theta(s^{1:T} | \mathcal{O}^{1:T}, \mathbf{u}^{0:T}, \mathbf{z}^{1:T})} \log \left[\frac{p(s^{1:T}, \mathcal{O}^{0:T}, \mathbf{u}^{0:T}, \mathbf{z}^{1:T})}{q_\theta(s^{1:T} | \mathcal{O}^{0:T}, \mathbf{u}^{0:T}, \mathbf{z}^{1:T})} \right] \\ &= \sum_{t=1}^T \left\{ r(s^t, \mathbf{u}^t) + \sum_{a=1}^n \log [p(u_a^t)] + \sum_{a=1}^n \log [p(z_a^t | s^t)] \right\} \end{aligned}$$

$$- \mathcal{D}_{KL} \left[q_{\theta} \left(s^t \mid s^{t-1}, \mathbf{u}^{t-1}, \mathbf{z}^t \right) \parallel p \left(s^t \mid s^{t-1}, \mathbf{u}^{t-1} \right) \right], \quad (3)$$

where $s^{1:T} \sim q_{\theta} \left(s^{1:T} \mid \mathcal{O}^{1:T}, \mathbf{u}^{0:T}, \mathbf{z}^{1:T} \right)$. Equation (3) holds under the condition that the action priors and observations of all agents are conditionally independent of each other given the state:

$$p(\mathbf{u}^t) = \prod_{a=1}^n p(u_a^t),$$

$$p(\mathbf{z}^t \mid s^t) = \prod_{a=1}^n p(z_a^t \mid s^t).$$

In order to maximize the evidence lower bound obtained, it is necessary to analyze each item in Equation (3). The first term $r(s^t, \mathbf{u}^t)$ can be maximised by reinforcement learning. $p(u_a^t)$ is an action prior, which is a constant. $p(z_a^t \mid s^t)$ represents that the latent state needs to be able to generate local observations of each agent. The negative KL divergence means we need to reduce the KL distance between the inference function $q_{\theta} \left(s^t \mid s^{t-1}, \mathbf{u}^{t-1}, \mathbf{z}^t \right)$ and the prior $p \left(s^t \mid s^{t-1}, \mathbf{u}^{t-1} \right)$. Since $p \left(s^t \mid s^{t-1}, \mathbf{u}^{t-1} \right)$ is unavailable, we adopt the method similar to SVQN by introducing an additional generative function $p_{\theta}^{\text{prior}} \left(s^t \mid s^{t-1}, \mathbf{u}^{t-1} \right)$ to approximate the prior $p \left(s^t \mid s^{t-1}, \mathbf{u}^{t-1} \right)$. In the next section, we will construct the specific neural network framework of SIDE to optimize the ELBO.

4.2 Framework

SIDE is based on the value decomposition methods, so the neural network architecture that implements the part of $q_{\pi}(\cdot)$ is consistent with the popular value decomposition methods. Each agent corresponds to a DRQN [8] as the agent network, whose input is the local observation of each agent and outputs the individual action-value Q_a . Several hypernetworks generate the parameters of the mixing network according to the state s . Finally, the mixing network merges all individual action-values to obtain a global action-value Q_{tot} . Through the above architecture diagram, $q_{\pi}(\cdot)$ is optimized by the value-based reinforcement learning methods. The only difference is that when the global state is not available, the parameters of the mixing network are determined by the latent state \tilde{s} derived from past information.

Next, we will introduce the two functions $p_{\theta}^{\text{prior}} \left(s^t \mid s^{t-1}, \mathbf{u}^{t-1} \right)$ and $q_{\theta} \left(s^t \mid s^{t-1}, \mathbf{u}^{t-1}, \mathbf{z}^t \right)$. First of all, to approximate the unknown prior distribution $p \left(s^t \mid s^{t-1}, \mathbf{u}^{t-1} \right)$, we need to construct a fitting function $p_{\theta}^{\text{prior}} \left(s^t \mid s^{t-1}, \mathbf{u}^{t-1} \right)$, here we use a standard VAE to implement. The VAE deduces the prior distribution of the current state s^t according to the past state s^{t-1} and the past actions \mathbf{u}^{t-1} of all agents. However, since the real state s^{t-1} is unavailable, we use the inferred latent state \tilde{s}^{t-1} . Besides, if the agent's action takes the form of one-hot in the discrete case, then \mathbf{u}^{t-1} will be very sparse, especially in the case of a large number of agents. Therefore, we design an embedding layer in the encoder and build a separate Multi-Layer Perceptron (MLP) for each agent in the decoder to output the reconstructed action to alleviate this problem. We think it can accelerate the training. Through the reparameterization trick, the prior distribution of the latent state $p_{\theta}^{\text{prior}} \left(\tilde{s}^t \mid \tilde{s}^{t-1}, \mathbf{u}^{t-1} \right)$ is easy to obtain.

Similarly, the state inference function $q_{\theta}(\cdot)$ is also constructed through generative models. In order to reduce the computational complexity, we have done two aspects of work. On the one hand, we regard the hidden output h_a^t of the recurrent neural network in the agent network corresponding to each agent as the integration of all past information of the corresponding individual agent, and assume that the hidden outputs of all agents \mathbf{h}^t contains all past information of the environment. Then the state inference function $q_{\theta} \left(s^t \mid s^{t-1}, \mathbf{u}^{t-1}, \mathbf{z}^{t-1} \right)$ can be rewritten as $q_{\theta} \left(s^t \mid \mathbf{h}^t \right)$. It should be noted that in order to stabilize the training process, we use the hidden outputs of the recurrent neural network generated by the target agent networks whose updates are relatively slow. On the other hand, instead of using a fully connected VAE, we use GALA, a variational graph auto-encoder with a symmetrical pair of the encoder and the decoder. One advantage of the graph network is that no matter how the number of agents changes, as long as the feature dimension of the node does not change, then the number of parameters that need to be trained will not change. Another advantage of introducing graph networks is that the relationship between nodes is taken into consideration. The relationship between nodes is represented by the adjacency matrix A . Last but not least, the inferred hidden state \tilde{s}^t is derived directly from the mean μ_{θ} of the distribution $q_{\theta} \left(s^t \mid \mathbf{h}^t \right)$, instead of sampling from $q_{\theta} \left(s^t \mid \mathbf{h}^t \right)$. That is, the hidden state \tilde{s}^t is obtained by concatenating the latent mean values of all nodes.

The framework of SIDE is shown in Figure 2. It is worth mentioning that the mixing network in SIDE can use different forms, so SIDE can be extended to various value decomposition methods.

4.3 Loss Function

In the above framework, a reinforcement learning process and two generative model optimization processes are included. To achieve the two purposes of optimal control planning and global state inference at the same time, the loss functions of the above two processes can be jointly optimized.

For the reinforcement learning process, we use the loss function of QMIX in this paper for simplicity, even though SIDE can be applied to various value decomposition multi-agent reinforcement learning algorithms. We express the parameters of the policy function $p_{\pi}(\cdot)$ as θ_R . The loss function of the reinforcement learning process can be computed by the following equation:

$$\mathcal{L}_{TD}(\theta_R) = (y_{tot} - Q_{tot}(\boldsymbol{\tau}, \mathbf{u} \mid \theta_R))^2,$$

where y_{tot} is the target joint value function and $y_{tot} = r + \gamma \max_{\mathbf{u}'} Q_{tot}(\boldsymbol{\tau}', \mathbf{u}' \mid \theta_R)$. θ_R is the parameters of the target network.

For prior model $p_{\theta_I}^{\text{prior}} \left(\tilde{s}^t \mid \tilde{s}^{t-1}, \mathbf{u}^{t-1} \right)$, we utilize the standard VAE training process, that is, to minimize the KL divergence between distribution $p_{\theta_I}^{\text{prior}} \left(\tilde{s}^t \mid \tilde{s}^{t-1}, \mathbf{u}^{t-1} \right)$ and $p \left(s^t \right)$, and the difference between the past information $\left(\left(\tilde{s}^{t-1} \right)', \left(\mathbf{u}^{t-1} \right)' \right)$ obtained by reconstruction and the actual past information $\left(\tilde{s}^{t-1}, \mathbf{u}^{t-1} \right)$. The two loss functions can be expressed as follows:

$$\mathcal{L}_{KL}^{\text{prior}}(\theta_I) = -\mathcal{D}_{KL} \left[p_{\theta_I}^{\text{prior}} \left(\tilde{s}^t \mid \tilde{s}^{t-1}, \mathbf{u}^{t-1} \right) \parallel p \left(s^t \right) \right],$$

$$\mathcal{L}_{RECON}^{\text{prior}}(\theta_I) =$$

$$\text{MSE} \left(\tilde{s}^{t-1}, \left(\tilde{s}^{t-1} \right)' \right) + \text{CROSS_ENTROPY} \left(\mathbf{u}^{t-1}, \left(\mathbf{u}^{t-1} \right)' \right),$$

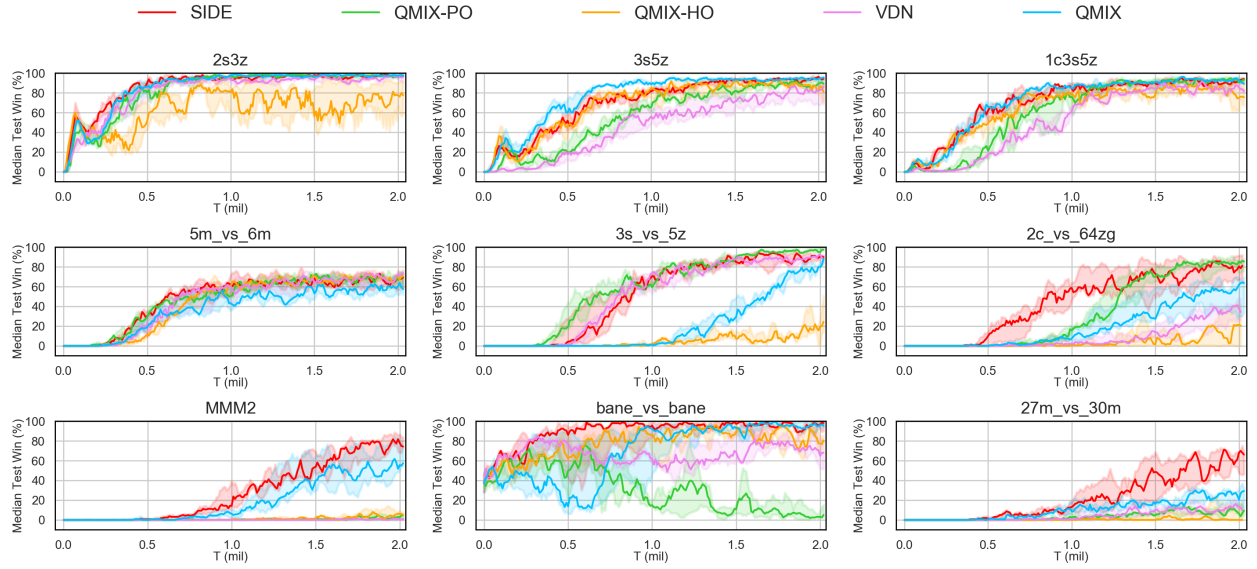


Figure 3: Overall results in different scenarios. Note that none of the algorithms except QMIX can get the true global state.

where $p(s^t)$ is the standard Gaussian distribution $\mathcal{N}(0, I)$. $MSE(\cdot)$ and $CROSS_ENTROPY(\cdot)$ represent the mean square error function and the cross entropy loss function respectively.

Finally, on the basis of the loss functions of GALA, The two equations that describes the training loss function of the state inference model $q_{\theta_I}(\tilde{s}^t | \mathbf{h}^t)$ are as follows:

$$\begin{aligned} \mathcal{L}_{KL}(\theta_I) &= -\mathcal{D}_{KL} \left[q_{\theta_I}(\tilde{s}^t | \mathbf{h}^t) \parallel p_{\theta_I}^{\text{prior}}(\tilde{s}^t | \tilde{s}^{t-1}, \mathbf{u}^{t-1}) \right], \\ \mathcal{L}_{RECON}(\theta_I) &= MSE(\mathbf{h}^t, (\mathbf{h}^t)'), \end{aligned}$$

where $(\mathbf{h}^t)'$ is the reconstructed hidden intermediate variables of agent networks and θ_I is the parameters of the inferred neural networks.

To sum up, we can get the loss function of the whole framework:

$$\mathcal{L}(\theta) = \mathcal{L}_{TD} + \mathcal{L}_{KL}^{\text{prior}} + \mathcal{L}_{RECON}^{\text{prior}} + \mathcal{L}_{KL} + \mathcal{L}_{RECON}.$$

By jointly optimizing the loss function, SIDE can infer the latent state that is most conducive to maximizing returns when the global state is unavailable, and gradually achieve optimal control.

5 EXPERIMENTS

We evaluate our proposed algorithm SIDE on the SMAC [35] platform based on StarCraft II. SMAC is a multi-agent testbed dedicated to solving Dec-POMDP problems. Different scenarios correspond to different problems, including heterogeneity, large action spaces, or a large number of agents. Although each agent can only obtain corresponding local observations, the global state can still be obtained by calling the interface of the SMAC. Many value decomposition methods use the global state for training. Even though this does not violate the paradigm of centralized training with decentralized execution, the global state is not available in some environments.

In this section, we first compare the performance of the three algorithms VDN, QMIX, and SIDE when the state information is not

given, and then compare the results with the performance of QMIX in the environment where global state information is available. In this way, it is judged whether SIDE can reconstruct valid state information based on past information.

Furthermore, we carry out the ablation studies to understand the contribution of all components of SIDE. In this paper, we focus on the contribution of the variational graph auto-encoders and the additional prior function. We define the latent state distribution dimension in the variational graph auto-encoder, also called the latent dimension, as N . To test whether the latent dimension has a significant influence on the performance of SIDE, we change the value of N in some scenarios and compare the experimental results.

5.1 Settings

The implementations of all algorithms are based on Pymarl, a multi-agent value decomposition algorithm integration platform. Both VDN and QMIX use the default hyperparameter settings. In addition, under the assumption that the perfect state information cannot be obtained, we concatenate the local observations or the hidden intermediate variables of all our own agents and replace the real state with this. We call these two variants QMIX-PO and QMIX-HO, respectively. For ease of comparison between the other baselines, SIDE is implemented on the basis of QMIX. The input adjacency matrix A of the variational graph auto-encoder is set to:

$$A_{ij} = \begin{cases} 1, & \text{if } i \text{ is alive and } j \text{ is alive} \\ 0, & \text{otherwise} \end{cases}, \quad \forall A_{ij} \in A.$$

VDN, QMIX-PO, QMIX-HO and SIDE have the same amount of information. By comparing their performance, we can prove that SIDE has an efficient solution to state inference. The above algorithms are also compared with vanilla QMIX to analyze the impact of the lack of global state information.

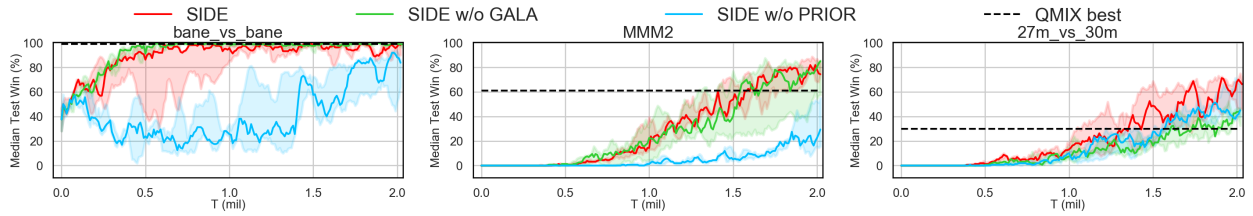


Figure 4: Win rates for SIDE and ablations in three representative scenarios.

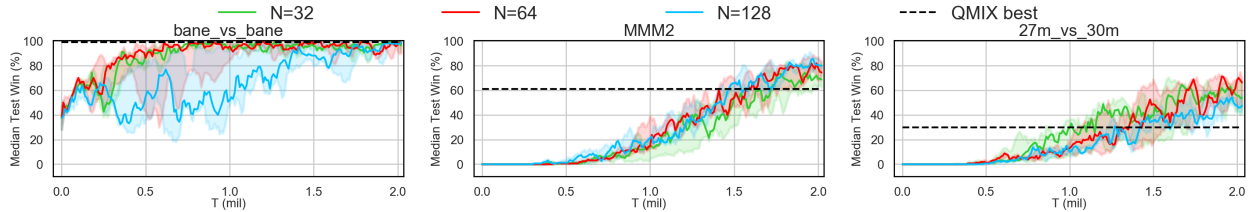


Figure 5: Influence of the different latent dimensions N for SIDE.

For ablation studies, we propose two alterations of SIDE. The first ablation replaces the graph auto-encoder in the original SIDE with the fully connected auto-encoder to test the importance of considering the relationship information between agents. And the second one is to remove the additional prior function $p_{\theta}^{\text{prior}}(\cdot)$ in the SIDE and directly use only the graph auto-encoder $q_{\theta}(\cdot)$ to infer the state. To sum up, we change the above two important components of SIDE while keeping other parts unchanged, and test the obtained two ablations in three representative scenarios which are hard or super hard. What’s more, we also carry out experiments to test the influence of different latent dimensions N .

In this paper, each algorithm runs five experiments independently with different random seeds to avoid the influence of outliers and we evaluate the algorithms every 10,000 time steps. The version of StarCraft II is 4.6.2 (B69232). In our experiments, the versions of GPU and CPU are Nvidia GeForce RTX 3090 and Intel(R) Xeon(R) Platinum 8280 respectively.

5.2 Validation

We get the results of the above five algorithms through experiments. In Figure 3, the performances of SIDE and the other four algorithms in different scenes are presented. The solid line represents the median win rate, and the 25-75% percentiles are shown. First, we observe the impact of the lack of information, that is, compare the results of the algorithms QMIX, QMIX-PO and QMIX-HO. In some scenarios, such as *3s5z*, *1c3s5z*, *MMM2*, *bane_vs_bane*, and *27m_vs_30m*, QMIX-PO performs significantly worse than QMIX. In addition, in the two scenes *5m_vs_6m* and *3s_vs_5z*, the performance of QMIX-PO and VDN are similar, and both are better than QMIX. This is in line with our intuition. For QMIX-HO, it performs well only in a few scenarios, including *3s5z*, *1c3s5z* and *bane_vs_bane*. In the *2c_vs_64zg* scenario, the result is more unexpected, because QMIX-PO is far better than QMIX. We think this is because there are too many enemy units, and only focusing on the

information of your units significantly improves the learning speed. The above results show that the definitions of state representations that are beneficial to learning in different scenarios are different: in some scenarios agents need global information, while in some scenarios agents tend to focus only on their own information.

From the results, we can easily see that the performance of SIDE in most scenarios close to the best one of the above four algorithms. It is also worth mentioning that SIDE performs far better than other algorithms in some super hard scenarios. This indicates that in different scenarios, SIDE can infer and reconstruct the latent state that is most helpful to the reinforcement learning process, regardless of whether the best state representation contains only its own information, or contains unobserved information, or even other unknown forms. Therefore SIDE does not need the true global state at all. Table 1 illustrates the median test win rate of different algorithms. The best performances of the above algorithms are bold and the second-best ones are underlined. Besides, we also compare SIDE with current state-of-the-art value decomposition methods and the results can be found in the appendix. In some complex scenarios, SIDE even outperforms all other popular baselines.

As can be seen from Figure 4, the information of the relationship between agents plays an important role in some scenarios such like *MMM2* and *27m_vs_30m*, but is redundant in few maps like *bane_vs_bane* where group relations are not valued. On the whole, we can infer that the graph auto-encoder has contributed to the outstanding performance of SIDE. As the another important component, the prior function significantly improve the learning speed of SIDE, which can be known from the poor performance of SIDE without the prior function. Last, we discuss the influence of different values of N . As depicted in Figure 5, all choices of N in the figure can typically generate satisfactory results, which means they all significantly outperform QMIX-PO and QMIX-HO. In practice, we set the latent dimension N to 64.

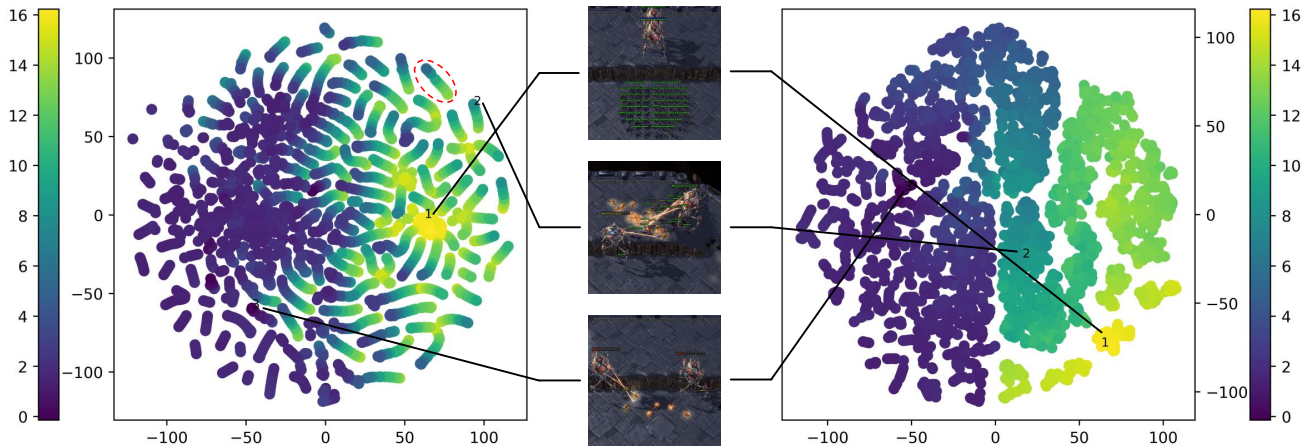


Figure 6: The 2D t-SNE embedding of true states on left, and latent states learned with SIDE on right. Each point is colored according to the state value. Three points corresponding to three states in different episodes are selected as examples. We find that the points in the red dashed circle almost belong to the state set of the same episode.

Table 1: Median performance of the test win ratio (%) in different scenarios.

scenario	SIDE	QMIX-PO	QMIX-HO	VDN	QMIX
2s3z	99	<u>98</u>	81	97	<u>98</u>
3s5z	96	90	93	86	<u>95</u>
1c3s5z	<u>95</u>	<u>95</u>	90	91	96
5m_vs_6m	<u>73</u>	<u>73</u>	74	74	63
3s_vs_5z	<u>94</u>	97	24	92	88
2c_vs_64zg	<u>85</u>	86	21	41	64
MMM2	81	5	8	1	<u>61</u>
bane_vs_bane	99	26	<u>96</u>	81	99
27m_vs_30m	71	13	13	16	<u>30</u>

5.3 Visualization

To visualize the latent states learned with SIDE, the t-SNE [24] plots are given in Figure 6. We compare the embedding of real states with that of latent states on *2c_vs_64zg* scenario after training has finished. For the real states, neighboring points in the embedding space often only indicate that the states they represent may belong to the same multi-agent decision-making process such as the points in the red circle, although their state values may be very different. In other words, there is no intuitive relationship between real states and state values, so it is not conducive to learning. On the contrary, points with the similar state values tend to cluster together in the embedding of latent states learned with SIDE. From the visualization, we can draw the conclusion that SIDE can infer the latent states that can accelerate the reinforcement learning process and we also understand why SIDE still has excellent performance even if global states are not available on *2c_vs_64zg* scenario.

6 CONCLUSION AND FUTURE WORK

In this paper, in order to infer the current state based on the past local observation, we propose SIDE, a novel algorithm that combines

multi-agent value decomposition and variational inference, so that no real state information is required during the entire training and execution process. Meanwhile, by jointly optimizing the two tasks of variational inference and optimal control, SIDE can promote the reconstruction of the state that is conducive to reinforcement learning, so that it far exceeds QMIX-PO, QMIX-HO and VDN in many tasks of SMAC, and even exceeds vanilla QMIX and other state-of-the-art methods.

SIDE can be easily extended to any multi-agent value decomposition algorithm with a mixing network, and the performances will be submitted in our future work. In addition, finding the interpretation of the reconstructed state is also a valuable research field, which will help enhance the robustness of multi-agent decision-making and help humans discover some unknown knowledge. It will also be the focus of our future work.

REFERENCES

- [1] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. 2016. Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, Daniel D. Lee, Masashi Sugiyama, Ulrike von Luxburg, Isabelle Guyon, and Roman Garnett (Eds.), 3837–3845. <https://proceedings.neurips.cc/paper/2016/hash/04df4d434d481c5bb723be1b6df1ee65-Abstract.html>
- [2] M. Egorov. 2015. Deep Reinforcement Learning with POMDPs.
- [3] Aleksandra Faust, Oscar Ramirez, Marek Fiser, Kenneth Oslund, Anthony G. Francis, James Davidson, and L. Tapia. 2018. PRM-RL: Long-range Robotic Navigation Tasks by Combining Reinforcement Learning and Sampling-Based Planning. *2018 IEEE International Conference on Robotics and Automation (ICRA)* (2018), 5113–5120.
- [4] Jakob N. Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. 2018. Counterfactual Multi-Agent Policy Gradients. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, Sheila A. McIlraith and Kilian Q. Weinberger (Eds.). AAAI Press, 2974–2982. <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/17193>
- [5] David Ha, Andrew M. Dai, and Quoc V. Le. 2017. HyperNetworks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net. <https://openreview.net/forum?id=rkpACe1lx>

- [6] Tuomas Haarnoja, Haoran Tang, P. Abbeel, and Sergey Levine. 2017. Reinforcement Learning with Deep Energy-Based Policies. In *ICML*.
- [7] Dongqi Han, Kenji Doya, and Jun Tani. 2020. Variational Recurrent Models for Solving Partially Observable Control Tasks. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net. <https://openreview.net/forum?id=r1LL4a4tDB>
- [8] M. Hausknecht and P. Stone. 2015. Deep Recurrent Q-Learning for Partially Observable MDPs. In *AAAI Fall Symposia*.
- [9] Shiyu Huang, Hang Su, Jun Zhu, and Ting Chen. 2020. SVQN: Sequential Variational Soft Q-Learning Networks. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net. <https://openreview.net/forum?id=r1xPh2VtPB>
- [10] Maximilian Igl, Luisa M. Zintgraf, Tuan Anh Le, Frank Wood, and Shimon Whiteson. 2018. Deep Variational Reinforcement Learning for POMDPs. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018 (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer G. Dy and Andreas Krause (Eds.). PMLR, 2122–2131. <http://proceedings.mlr.press/v80/igl18a.html>
- [11] Shariq Iqbal and Fei Sha. 2019. Actor-Attention-Critic for Multi-Agent Reinforcement Learning. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA (Proceedings of Machine Learning Research, Vol. 97)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.). PMLR, 2961–2970. <http://proceedings.mlr.press/v97/iqbal19a.html>
- [12] Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Z. Leibo, David Silver, and Koray Kavukcuoglu. 2017. Reinforcement Learning with Unsupervised Auxiliary Tasks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net. <https://openreview.net/forum?id=SJ6yPD5xg>
- [13] Rico Jonschkowski, Divyam Rastogi, and Oliver Brock. 2018. Differentiable Particle Filters: End-to-End Learning with Algorithmic Priors. (2018).
- [14] Peter Karkus, David Hsu, and Wee Sun Lee. 2018. Particle Filter Networks with Application to Visual Localization. In *CoRL*.
- [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. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net. <https://openreview.net/forum?id=SJxu5iR9KQ>
- [16] Diederik P. Kingma and Max Welling. 2014. Auto-Encoding Variational Bayes. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.). <http://arxiv.org/abs/1312.6114>
- [17] Thomas Kipf and M. Welling. 2016. Variational Graph Auto-Encoders. *ArXiv abs/1611.07308* (2016).
- [18] D. Koller and N. Friedman. 2009. Probabilistic Graphical Models - Principles and Techniques.
- [19] Guillaume Lample and Devendra Singh Chaplot. 2017. Playing FPS Games with Deep Reinforcement Learning. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, Satinder P. Singh and Shaul Markovitch (Eds.). AAAI Press, 2140–2146. <http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14456>
- [20] P. Larrañaga. 2002. An Introduction to Probabilistic Graphical Models. In *Estimation of Distribution Algorithms*.
- [21] Sergey Levine. 2018. Reinforcement Learning and Control as Probabilistic Inference: Tutorial and Review. *ArXiv abs/1805.00909* (2018).
- [22] Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. 2017. Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (Eds.). 6379–6390. <https://proceedings.neurips.cc/paper/2017/hash/68a9750337a418a86fe06c1991a1d64c-Abstract.html>
- [23] Xiao Ma, Peter Karkus, David Hsu, Wee Sun Lee, and Nan Ye. 2020. Discriminative Particle Filter Reinforcement Learning for Complex Partial Observations. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net. https://openreview.net/forum?id=HJl8_eHYvS
- [24] L. V. D. Maaten and Geoffrey E. Hinton. 2008. Visualizing Data using t-SNE. *Journal of Machine Learning Research* 9 (2008), 2579–2605.
- [25] Anuj Mahajan, Tabish Rashid, Mikayel Samvelyan, and Shimon Whiteson. 2019. MAVEN: Multi-Agent Variational Exploration. 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.). 7611–7622. <https://proceedings.neurips.cc/paper/2019/hash/f816dc0acf67498e10496222e9db10-Abstract.html>
- [26] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, Ioannis Antonoglou, Daan Wierstra, and Martin A. Riedmiller. 2013. Playing Atari with Deep Reinforcement Learning. *ArXiv abs/1312.5602* (2013).
- [27] Pol Moreno, Edward Hughes, Kevin R. McKee, Bernardo Ávila Pires, and Théophane Weber. 2021. Neural Recursive Belief States in Multi-Agent Reinforcement Learning. *ArXiv abs/2102.02274* (2021).
- [28] F. Oliehoek and Chris Amato. 2016. A Concise Introduction to Decentralized POMDPs. In *SpringerBriefs in Intelligent Systems*.
- [29] OpenAI, I. Akkaya, Marcin Andrychowicz, Maciej Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, J. Schneider, N. Tezak, Jerry Tworek, P. Welinder, Lilian Weng, Qiming Yuan, W. Zaremba, and Lei Zhang. 2019. Solving Rubik’s Cube with a Robot Hand. *ArXiv abs/1910.07113* (2019).
- [30] Shirui Pan, Ruiqi Hu, Guodong Long, Jing Jiang, Lina Yao, and Chengqi Zhang. 2018. Adversarially Regularized Graph Autoencoder for Graph Embedding. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, Jérôme Lang (Ed.). ijcai.org, 2609–2615. <https://doi.org/10.24963/ijcai.2018/362>
- [31] Jiwoong Park, Minsik Lee, Hyung Jin Chang, Kyuewang Lee, and Jin Young Choi. 2019. Symmetric Graph Convolutional Autoencoder for Unsupervised Graph Representation Learning. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*. IEEE, 6518–6527. <https://doi.org/10.1109/ICCV.2019.00662>
- [32] Peng Peng, Ying Wen, Y. Yang, Quan Yuan, Zhenkun Tang, Haitao Long, and J. Wang. 2017. Multiagent Bidirectionally-Coordinated Nets: Emergence of Human-level Coordination in Learning to Play StarCraft Combat Games. *ArXiv: Artificial Intelligence* (2017).
- [33] Tabish Rashid, Mikayel Samvelyan, Christian Schröder de Witt, Gregory Farquhar, Jakob N. Foerster, and Shimon Whiteson. 2018. QMIX: Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018 (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer G. Dy and Andreas Krause (Eds.). PMLR, 4292–4301. <http://proceedings.mlr.press/v80/rashid18a.html>
- [34] Guillaume Salha, Romain Hennequin, and M. Vazirgiannis. 2019. Keep It Simple: Graph Autoencoders Without Graph Convolutional Networks. *ArXiv abs/1910.00942* (2019).
- [35] Mikayel Samvelyan, Tabish Rashid, C. S. D. Witt, Gregory Farquhar, Nantas Nardelli, Tim G. J. Rudner, Chia-Man Hung, P. Torr, Jakob N. Foerster, and S. Whiteson. 2019. The StarCraft Multi-Agent Challenge. In *AAMAS*.
- [36] Kyunghwan Son, Daewoo Kim, Wan Ju Kang, David Hostallero, and Yung Yi. 2019. QTRAN: Learning to Factorize with Transformation for Cooperative Multi-Agent Reinforcement Learning. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA (Proceedings of Machine Learning Research, Vol. 97)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.). PMLR, 5887–5896. <http://proceedings.mlr.press/v97/son19a.html>
- [37] Sainbayar Sukhbaatar, Arthur Szlam, and Rob Fergus. 2016. Learning Multi-agent Communication with Backpropagation. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, Daniel D. Lee, Masashi Sugiyama, Ulrike von Luxburg, Isabelle Guyon, and Roman Garnett (Eds.). 2244–2252. <https://proceedings.neurips.cc/paper/2016/hash/55b1927fdafef39c48e5b73b5d61ea60-Abstract.html>
- [38] Peter Sunehag, G. Lever, A. Gruslys, Wojciech Czarnecki, V. Zambaldi, Max Jaderberg, Marc Lanctot, Nicolas Sonnerat, Joel Z. Leibo, K. Tuyls, and T. Graepel. 2018. Value-Decomposition Networks For Cooperative Multi-Agent Learning. *ArXiv abs/1706.05296* (2018).
- [39] J. Wang, Zhizhou Ren, T. Liu, Yang Yu, and C. Zhang. 2020. QPLEX: Duplex Dueling Multi-Agent Q-Learning. *ArXiv abs/2008.01062* (2020).
- [40] Rose E. Wang, M. Everett, and J. How. 2020. R-MADDPG for Partially Observable Environments and Limited Communication. *ArXiv abs/2002.06684* (2020).
- [41] Tonghan Wang, Heng Dong, Victor R. Lesser, and Chongjie Zhang. 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 (Proceedings of Machine Learning Research, Vol. 119)*. PMLR, 9876–9886. <http://proceedings.mlr.press/v119/wang20f.html>
- [42] Tonghan Wang, Tarun Gupta, Anuj Mahajan, Bei Peng, S. Whiteson, and C. Zhang. 2020. RODE: Learning Roles to Decompose Multi-Agent Tasks. *ArXiv abs/2010.01523* (2020).
- [43] Ying Wen, Yaodong Yang, Rui Luo, Jun Wang, and Wei Pan. 2019. Probabilistic Recursive Reasoning for Multi-Agent Reinforcement Learning. *ArXiv abs/1901.09207* (2019).
- [44] Pengfei Zhu, X. Li, and P. Poupart. 2017. On Improving Deep Reinforcement Learning for POMDPs. *ArXiv abs/1804.06309* (2017).