

GCS: Graph-Based Coordination Strategy for Multi-Agent Reinforcement Learning

Jingqing Ruan¹², Yali Du^{*3}, Xuantang Xiong¹, Dengpeng Xing¹, Xiyun Li¹, Linghui Meng¹,
Haifeng Zhang¹, Jun Wang⁴, Bo Xu^{*1}

¹Institute of Automation, Chinese Academy of Sciences, Beijing, China

²School of Future Technology, University of Chinese Academy of Sciences, Beijing, China

³King’s College London, ⁴University College London

{ruanjingqing2019,xiongxuantang2021,lixiyun2020,dengpeng.xing,menglinghui2019,haifeng.zhang,xubo}@ia.ac.cn
yali.du@kcl.ac.uk,jun.wang@cs.ucl.ac.uk

ABSTRACT

Many real-world scenarios involve a team of agents that have to coordinate their policies to achieve a shared goal. Previous studies mainly focus on decentralized control to maximize a common reward and barely consider the coordination among control policies, which is critical in dynamic and complicated environments. In this work, we propose factorizing the joint team policy into a graph generator and graph-based coordinated policy to enable coordinated behaviours among agents. The graph generator adopts an encoder-decoder framework that outputs directed acyclic graphs (DAGs) to capture the underlying dynamic decision structure. We also apply the DAGness-constrained and DAG depth-constrained optimization in the graph generator to balance efficiency and performance. The graph-based coordinated policy exploits the generated decision structure. The graph generator and coordinated policy are trained simultaneously to maximize the discounted return. Empirical evaluations on Collaborative Gaussian Squeeze, Cooperative Navigation, and Google Research Football demonstrate the superiority of the proposed method. The code is available at https://github.com/Amanda-1997/GCS_aamas337.

KEYWORDS

Action Coordination Graph; Multi-Agent Systems; Reinforcement Learning

ACM Reference Format:

Jingqing Ruan¹², Yali Du^{*3}, Xuantang Xiong¹, Dengpeng Xing¹, Xiyun Li¹, Linghui Meng¹, Haifeng Zhang¹, Jun Wang⁴, Bo Xu^{*1}. 2022. GCS: Graph-Based Coordination Strategy for Multi-Agent Reinforcement Learning. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Online, May 9–13, 2022*, IFAAMAS, 9 pages.

1 INTRODUCTION

Multi-agent reinforcement learning (MARL) has shown exceptional results in many real-life applications, such as multiplayer games [1, 2], traffic control [3], and social dilemmas [4]. A suitable control policy is extremely important in multi-agent systems (MASs). One choice is to treat the MAS as a single agent and adopt a centralized control policy [5, 6]; however, this approach is constrained by poor

scalability for high-dimensional state and action spaces. On the contrary, decentralized control [2, 7–10] allows agents to make decisions independently, but struggles to enable coordinated behaviors on complex tasks. Taking traffic flow as an example, when multiple vehicles are trying to cross an intersection without traffic lights, most likely, the traffic will become congested if all vehicles take actions simultaneously without a rational sequence. This problem may be solved, however, if those vehicles move in an orderly way based on some coordination structure. This example shows that it is imperative to improve on a fully decentralized decision-making process, and a solution to alleviate the above issue is to develop a coordinated control policy to obtain cooperative behaviors.

Several approaches have been reported to address the problem of action coordination. BiAC [11] mainly focuses on coordination of the asynchronous decisions of two agents. The multi-agent rollout algorithm [12] provides a theoretical view of executing a local rollout with some coordinating information, but is limited to an agent-by-agent decision dependency structure. Although these works investigate the action execution order of two-agent and multi-agent systems, they are still insufficient to characterize the complicated and dynamic underlying decision dependency structure of a general multi-agent system. Moreover, we assert that representing the underlying decision dependency structure and using this to control the action execution is essential to improving coordination.

In this work, we propose a graph-based coordination strategy (GCS) that learns coordinated behaviors through factorizing the joint team policy into a graph generator and a graph-based coordinated policy. The former aims to learn an action coordination graph (ACG) that properly represents the decision dependency. The latter further coordinates the dependent behaviors among agents exploiting the underlying decision dependency. We train the graph generator and the graph-based coordinated policy simultaneously to maximize the discounted return. For the ACG we employ directed acyclic graphs (DAGs), whose nodes represent agents and whose directed edges denote action dependencies of the associated agents. Moreover, we propose using the DAGness-constrained and DAG depth-constrained optimization in the graph generator to balance efficiency and performance.

The contributions of this paper can be summarized as follows:

- As far as we know, we are the first to introduce directed acyclic graphs to action coordination, dynamically representing the underlying decision dependencies of MAS.

* Corresponding author.

- We propose a DAGness-constrained and DAG depth-constrained optimization in the graph generator, achieving a trade-off between decision-making efficiency and performance.
- Empirical evaluations on several challenging MARL benchmarks (Collaborative Gaussian Squeeze, Cooperative Navigation, and Google Football) show that our method can achieve superior performance and obtain meaningful results consistent with intuitive expectations.

2 RELATED WORK

Deep reinforcement learning has been successfully applied to addressing complex decision problems [13–17]. Due to the widespread existence of multi-agent tasks, MARL has attracted increasing attention, and learning appropriate control policies is important to obtain the maximum cumulative discounted return. Based on the structures of their execution schemes, we classify the existing approaches into three categories.

First, a fully independent execution scheme allows agents to determine actions according to their individual policies without any interaction. One line of research, such as IQL [18], VDN [7], QMIX [8], and QTRAN [19], focuses on value-based methods, which assign each agent an independent policy for execution. Another line of research, including MAAC [10], COMA [20], and LIIR [9], extends the actor-critic algorithm [21] to the multi-agent case, where each actor represents an individual policy for an agent.

Second, the communication-based independent execution scheme is widely used, which allows the use of extensive information in its individual decision making [22]. In this scheme, agents learn how to transmit informative messages and to process the messages during training. Then agents exchange the messages to determine their actions individually during independent execution. Representative methods [6, 23–27] autonomously learn communication protocols that are required in generating informative messages: these determine whom to communicate with and what messages to transmit for assisted decision making.

Third, the coordinated execution scheme, where agents develop their policies conditioned on other agents’ actions and make decisions in a coordinated manner, is important in MAS. There are some methods that implicitly model the coordinated behaviours from the perspective of a coordination graph. DCG [28] uses pairwise graphs to propagate beliefs for joint decisions, while DICG [29] focuses on generating a coordination graph with soft edge weights for message passing. DGN [30] uses a graph attention network as an embedding extractor to assist in the decision making. Furthermore, some methods have been proposed to explicitly model the coordinated behaviours in order, such as Bi-AC [11] and multi-agent rollout [12], which propose utilizing two-agent and agent-by-agent dependency structures, respectively, to help agents make decisions in order and promote action coordination. Similar but essentially distinct, we introduce a mechanism to learn the underlying DAG structure that represents the decision dependency among agents.

Moreover, the generation of the DAGs is an essential part of our work. Recently, some continuous optimization approaches [31–34] have been proposed to recover the DAGs through structure learning. The method [35], closely related to our work, uses reinforcement learning as its search strategy to maximize a predefined

score function. Borrowing from an idea of [35], which obtains the DAG using reinforcement learning, we construct a graph generator module to generate the DAG structure as an action coordination graph and regard the extrinsic reward as the incentives to jointly train with MARL tasks.

3 PROBLEM SETUP

MMDPs. Cooperative multi-agent problems can be modeled as multiagent Markov decision processes (MMDPs) [36], which can be expressed as a tuple $\langle \{\mathcal{I}\}, \mathcal{S}, \{\mathcal{U}^i\}_{i=1}^N, \mathbb{P}, r, \gamma \rangle$. $i \in \mathcal{I}$ is the i^{th} player, \mathcal{S} is the global state space, and \mathcal{U}^i denotes the action space for the i^{th} player. We label $\mathbf{u} := (u^1, \dots, u^N)$ the joint actions for all players. Intuitively, the agent i will select an individual action u^i to perform and execute it. \mathbb{P} is the transition dynamics, and $\mathbb{P}(\cdot|s, \mathbf{u})$ gives the distribution of the next state at state s taking action \mathbf{u} . All agents share the same reward function $r(s, \mathbf{u}) : s \times \mathbf{u} \rightarrow R$. $\gamma \in (0, 1)$ denotes a discount factor, and $\tau = (s_0, \mathbf{u}_0, s_1, \dots)$ denotes the trajectory induced by the policy $\pi = \{\pi^i\}_{i=1}^N$. All the agents coordinate together to maximize the cumulative discounted return $\mathbb{E}_{\tau \sim \pi} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, \mathbf{u}_t) \right]$.

Factored MMDPs. We formalize our problem based on MMDPs as factored MMDPs. Different from MMDPs, where all actions are taken simultaneously and do not depend on each other, we endow the hierarchy order to the joint action based on the learned DAG structure \mathcal{G} , called an action coordination graph (ACG). The adjacency matrix A representing the ACG denotes the decision dependency from the graph generator ρ . With A , we can define $\pi = \prod_{i=1}^n \pi^i(u^i|o^i, u^{pa(i)})$ as the graph-based coordinated policy for the i^{th} player, where o^i is the observation of the i -th player, $pa(i)$ are the parents of agent i , and $u^{pa(i)}$ are the actions taken by the parents, whose order is generated from A . Note that a fully decentralized policy is a special case of our graph-based coordinated policy where none of the agents have parents.

Figure 1 gives an example of the DAG and the relationships between nodes. The nodes in \mathcal{G} correspond to the agents in the MAS, and the parent-child relationships represent the hierarchical decision dependencies among agents. For example, the case that the node D in the graph has two parents $\{B, C\}$ illustrates that the action taken by the agent D is constrained by B and C .

Now, the graph-based coordination strategy is factorized as:

$$\pi(\mathbf{u}, A|s) = \rho(A|s) \prod_{i=1}^n \pi^i(u^i|o^i, u^{pa(i) \sim A}), \quad (1)$$

where ρ is the graph generator, and π^i is graph-based coordinated policy.

4 METHODOLOGY

The overall goal is to maximize the cumulative return, denoted as:

$$\eta = \mathbb{E}_{A \sim \rho, \mathbf{u} \sim \pi(\cdot|s, A)} \left[\sum_{k=0}^{\infty} \gamma^k r(s_{t+k}, \mathbf{u}_{t+k}) \right]. \quad (2)$$

Now we further elaborate on the derivation of the graph-based coordinated policy $\pi_{\theta_i}(u^i|o^i, u^{pa(i)})$ and the graph generator $\rho_{\phi}(A|s)$, respectively.




Figure 1: The left panel is a schematic diagram of a directed acyclic graph \mathcal{G} with vertices $|V| = 5$ and edges $|\mathcal{E}| = 7$. The right panel shows the parent-child relationships of \mathcal{G} .

4.1 Graph-based Coordinated Policy

Given a known graph generator ρ (elaborated in Section 4.2), we have $A \sim \rho$ to represent the underlying decision dependency. Based on it, we can denote the decision policy as $\pi = \prod_{i=1}^N \pi^i(u^i | o^i, u^{pa(i) \sim A})$, called graph-based coordinated policy. We explore graph-based coordinated policy π^* that obtains the final joint action $\{u^i\}_{i=1}^N$ as follows.

With Equations (1) and (2), we can write the expected return for agent i as:

$$\begin{aligned} \eta^i &= \mathbb{E}_{s \sim p^\pi, (u, A) \sim \pi} [\pi(u, A | s) Q_\pi^i(s, u)] \\ &= \mathbb{E}_{s \sim p^\pi, u^i \sim \pi^i, A \sim \rho} \left[\rho(A | s) \prod_{i=1}^N \pi^i(u^i | o^i, u^{pa(i) \sim A}) Q_\pi^i(s, u) \right]. \end{aligned} \quad (3)$$

The graph-based coordinated policy $\pi = \{\pi_1, \pi_2, \dots, \pi_N\}$ for N agents can be parameterized by $\theta = \{\theta_1, \theta_2, \dots, \theta_N\}$. Correspondingly, the gradient of the expected return for agent i is expressed as:

$$\begin{aligned} \nabla_{\theta^i} \eta(\theta^i) &= \nabla_{\theta^i} \eta^i \\ &= \mathbb{E}_{s \sim p^\pi, u^i \sim \pi^i} \left[\nabla_{\theta^i} \log \pi^i(u^i | o^i, u^{pa(i)}) \sum_A \rho(A | s) Q_\pi^i(s, u) \right]. \end{aligned} \quad (4)$$

By applying the mini-batch technique to the off-policy training, the gradient can be approximately estimated as:

$$\nabla_{\theta^i} \eta(\theta^i) = \mathbb{E}_{(s, o, A, u) \sim \mathcal{D}} \left[\nabla_{\theta^i} \log \pi^i(u^i | o^i, u^{pa(i)}) Q_\pi^i(s, u) \right], \quad (5)$$

where \mathcal{D} is the experience replay buffer, recording experiences of all agents. Moreover, the centralized action-value function Q_π^i can be updated as:

$$\mathcal{L}(\phi) = E_{(s, o, A, u) \sim \mathcal{D}} \left[\left(Q_\pi^i(s, u) - y^i \right)^2 \right] \quad (6)$$

where $y^i = r + \gamma \max_{u'} Q_{\pi^-}^i(s', u')$ is the learning target and $Q_{\pi^-}^i$ is the target network parameterized by θ^{i-} .

During the training process, the graph-based coordinated policy $\pi = \{\pi_1, \pi_2, \dots, \pi_N\}$ and the graph generator ρ are updated iteratively. We will describe how to find the graph generator $\rho(\cdot)$ under a given policy π .

4.2 Graph Generator

The graph generator aims to generate the DAG \mathcal{G} to define the decision dependency among agents. We will introduce it in detail from

three aspects: (a) DAGness constraint, (b) DAG Depth constraint, and (c) optimization objective.

DAGness constraint. The acyclicity constraint is an important issue in our problem setting. In this work, we also use the penalty terms like [31, 34, 35] to ensure acyclicity. The result in [31] shows that the directed graph \mathcal{G} with binary adjacency matrix A is acyclic if and only if:

$$g(A) := \text{trace}(e^{A \circ A}) - d = 0, \quad (7)$$

where $e^{A \circ A}$ is the matrix exponential, $A \circ A$ guarantees the non-negativity, and d is the number of nodes in the DAGs. The ‘trace’ of a matrix is defined as the sum of the diagonal elements [31]. The constraint function g should satisfy that: (a) its derivatives are computable, and (b) g can be the measurement of DAGs.

DAG Depth constraint. Moreover, taking the trade-off between efficiency and performance into account, we claim that the maximum depth of graph structure should be adjustable over different tasks. Therefore, we propose an alternative constraint to control the hierarchy of the generated DAGs as follows.

DEFINITION 1. A square matrix A is a Nilpotent Matrix [37], if

$$A^k = O \text{ and } A^{k-1} \neq O, \exists k \in \mathbb{Z}^+,$$

where O is the zero matrix and A is called the Nilpotent of index k .

PROPOSITION 1. Let A be an adjacency matrix for a directed acyclic graph, then the maximal length between any two nodes i and j is k if A is Nilpotent of index k .

We provide a detailed proof of proposition 1 in Appendix A.1. Here, we define $c(A^k) := \text{sum}(A^k) = \sum_i \sum_j A_{ij}^k = 0$, which is equivalent to $A^k = O$. We remark that as long as constraint $A^k = O$ holds, it can be guaranteed that the maximal length between any two nodes i and j of the DAG does not exceed k .

Optimization objective. Based on the foregoing, we can optimize ρ parameterized by φ with the maximal length k by:

$$\begin{aligned} \max \eta(\varphi) &= \mathbb{E} \left[\sum_{k=0}^{\infty} \gamma^k r^\rho(s_{t+k}, u_{t+k}) \right], \\ \text{s.t. } g(W) &= 0, c(W^k) = 0 \end{aligned} \quad (8)$$

where $W = \rho_\varphi(\cdot)$ denotes the weight matrix generated from the graph generator ρ . Then the weight matrix is modeled as a Bernoulli distribution, from which the binary adjacency matrix A is sampled. Here, we use the constraints of the weight matrix $g(W)$ and $c(W^k)$ to approximate those of the adjacency matrix $g(A)$ and $c(A^k)$ due to the consistency of representing the graph structure. With this approximation, we restate η^i as:

$$\eta^i = \mathbb{E} \left[\rho(W | s) \prod_{i=1}^N \pi^i(u^i | o^i, u^{pa(i)}) Q_\pi^i(s, u) \right]. \quad (9)$$

Fixing graph-based coordinated policy π , we approximate the graph generator ρ as follows. We augment the original problem shown in Equation (8) with a quadratic penalty using the augmented Lagrangian technique [38]:

$$\begin{aligned} \min_{\varphi} \quad & -\eta(\varphi) + \frac{\xi}{2} \left[\|g(W)\|^2 + \|c(W^k)\|^2 \right], \\ \text{s.t. } \quad & g(W) = 0, c(W^k) = 0 \end{aligned} \quad (10)$$

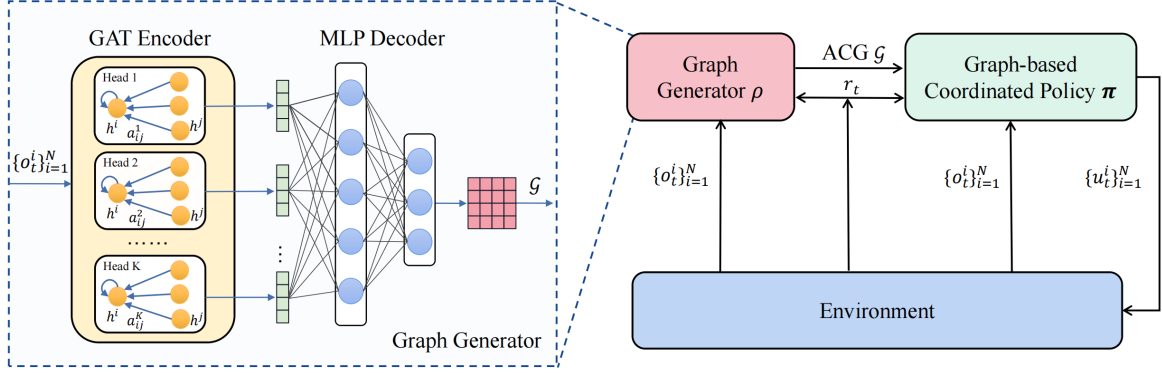


Figure 2: The proposed framework for the interaction process. The left is the graph generator including an encoder-decoder neural network model which is used to generate the DAG from the observed data.

with the penalty $\xi > 0$.

Next, we convert the Equation (10) to an unconstrained Lagrangian function:

$$L(\varphi, \lambda_1, \lambda_2) = -\eta(\varphi) + \frac{\xi}{2} \left[\|g(W)\|^2 + \|c(W^k)\|^2 \right] + \lambda_1 g(W) + \lambda_2 c(W^k). \quad (11)$$

PROPOSITION 2. *The gradient for Equation (11) to optimize the coordination graph generation policy can be derived as follows:*

$$\begin{aligned} \nabla_{\varphi} L(\varphi, \lambda_1, \lambda_2) = E_{s \sim p, W \sim \rho_{\varphi}(\cdot|s)} \left[\nabla_{\varphi} \log \rho_{\varphi}(W|s) \sum_{u^i} \pi^i(u^i|o^i, \right. \\ \left. u^{pa(i)}) \cdot Q^i(o^i, u^i) - \lambda_1 (e^{W^o W})^T \cdot 2W - \right. \\ \left. \lambda_2 \sum_{i,j} [k W^k W^{-1}]_{ij} \right], \end{aligned}$$

We provide a detailed proof of proposition 2 in Appendix A.2.

In proposition 2, we remark that after considering the influence of various decision dependencies on the reinforcement learning tasks, we can obtain the underlying graph structure that makes the best response to MARL tasks.

4.3 Implementation Details

As shown in Figure 2, the proposed framework include the graph-based coordinated policy and the graph generator, which will be elaborated below.

Graph-based Coordinated Policy. The graph-based coordinated policy can be obtained from the standard multi-agent actor-critic framework. As for the policy π^i of the actor, we use the RNN network with the stochastic policy gradient to model the action distributions. The critic used to criticize the joint actions made by the actors is a three-layer feed-forward neural network activated by the ReLU units, denoted as $f_{critic}(o_t, u_t) = \text{ReLU}(\text{MLP}(o_t, u_t))$.

Graph Generator. As shown in Figure 2, the graph generator $\rho_{\varphi}(\mathbf{o})$ adopt an encoder-decoder module used to find the ACQ \mathcal{G} . The GAT-based encoder can model the interplay of agents and extract the further latent representations. The MLP-based decoder is used

to recover the pairwise relationship between agents to generate an ACQ. The graph generator takes the local observations of the agents as input and outputs the ACQ \mathcal{G} to obtain the decision dependency for the decision-making process of the graph-based coordinated policy, elaborated as follows.

The graph generator contains two sub-modules. Firstly, we use the graph attention network (GAT) [39] as the attention-based encoder to extract the latent information for the graph structure generation. First, the simple feature is extracted by a multi-layer perceptron (MLP) as an initial step:

$$\{h_t^i\}_{i=1}^N = f_{MLP}(\{o_t^i\}_{i=1}^N, \{u_{t-1}^i\}_{i=1}^N). \quad (12)$$

Due to the sufficiently expressive power of GAT, we use it to extract the further latent information of the simple feature. We compute the importance coefficients through the attention mechanism:

$$a_t^{ij} = \frac{\exp(h_t^i \mathbf{W} (h_t^j \mathbf{W})^T / \sqrt{d})}{\sum_{k \in \{I\}^{-i}} \exp(h_t^i \mathbf{W} (h_t^k \mathbf{W})^T / \sqrt{d})}, \sum_{j \in \{I\}^{-i}} a_t^{ij} = 1, \quad (13)$$

where $\mathbf{W} \in R^{d' \times d}$ is a learnable weight matrix, d' and d denote the dimensions of the input vector and the latent vector, respectively, and $k \in \{I\}^{-i}$ indexes other agents except the agent i .

Then the multi-head attention is used to stabilize the learning process of self-attention, and the final latent feature is as follows:

$$h_t^{i'} = \sigma \left(\frac{1}{M} \sum_{m=1}^M \sum_{j \in \{I\}^{-i}} a_t^{ij} \mathbf{W}^m h_t^j \right). \quad (14)$$

where M is the number of attention heads. $a_{t,m}^{ij}$ are importance coefficients computed by the k -th attention mechanism, and \mathbf{W}^m is the corresponding weight matrix.

Another sub-module in the graph generator is the decoder that generates a weight matrix used to sample the graph structure. Since the GAT-based encoder has already provided sufficiently expressive features among agents, a single-layer decoder is enough to easily construct the pairwise relationship between the encoder outputs

to find a better structure of DAG for the decision policy.

$$f_t^{ij}(\mathbf{h}_t^l, \mathbf{h}_t^r, u) = \sigma(u^T \tanh(\mathbf{W}_l \mathbf{h}_t^l + \mathbf{W}_r \mathbf{h}_t^r)), \quad (15)$$

where \mathbf{h}_t^l and \mathbf{h}_t^r are agents' higher-level representations from two encoder outputs, $\mathbf{W}_l, \mathbf{W}_r \in \mathbb{R}^{d_h \times d_e}$, $u \in \mathbb{R}^{d_h \times 1}$ are trainable parameters, and d_h, d_e are the hidden dimension and encoder output dimension, respectively.

Moreover, the logistic sigmoid function $\sigma(\cdot)$ generates the probability for constructing the Bernoulli distribution from which the binary adjacency matrix A is sampled. The binary adjacency matrix forms a directed graph corresponding to the ACG \mathcal{G} . Here, we denote this graph generation process as $\mathcal{G} \sim \rho_\varphi(\mathbf{h})$.

Parameter Setting. In the graph generator, the attention head in the GAT encoder is set to 8, the stacked attentional layers are set to 4, and the hidden units in the MLP is set to 64. In the graph-based coordinated policy, the actor critic architecture is adopted. The recurrent layer comprised of a GRU with a 64-dimensional hidden state, with a fully-connected layer before and after, is used as the actor. The critic is a two-layer MLP with the ReLU activation.

4.4 Algorithm Description

The main procedures are summarized in Algorithm 1, where $\nabla_\varphi L(\varphi)$, $\eta_{\theta^i}(\theta^i)$, and $\nabla_\phi \mathcal{L}(\phi)$ are optimized. Our ultimate goal is to obtain the graph-based coordinated policy $\{\pi^i\}_{i=1}^n$. The graph generator ρ is an intermediate used to access an excellent graph structure in guiding the decision-making sequence among agents to achieve a high degree of multi-agent coordination. The graph generator is a pluggable module that can be replaced by other algorithms for solving the DAGs. Note that DAGs are necessary because we need an execution structure that can determine a clear sequence of the actions, and therefore it should be directed and not circular. The policy solver is a universal module, which, in general, one can choose from a diverse set of cooperative MARL algorithms [2, 10, 40].

5 EXPERIMENTS

We evaluate the effectiveness of our algorithm on three different environments: Collaborative Gaussian Squeeze¹ [19], Cooperation Navigation² [2], and Google Research Football³ [41].

5.1 Experimental Setting

Collaborative Gaussian Squeeze (CGS). As an extension of Multi-domain Gaussian Squeeze (MGS) [19], Collaborative Gaussian Squeeze is a challenging environment for evaluating coordination. In MGS, there exist K domains $[(\mu_1, \sigma_1), \dots, (\mu_K, \sigma_K)]$ in the system. The system contains N agents, and each agent i can take actions a_i within range of $\{-10, -9, \dots, 0, 1, \dots, 8, 9, 10\}$. The prior $s_i \in [0, 0.2]$ given by the environment represents the unit-level resource for each agent i . The total amount of resources mobilized by all agents is denoted as $f(\mathbf{u}) = \sum_i s_i \times a_i$. The goal is to maximize the joint reward $G(\mathbf{u}) = \sum_{k=1}^K f(\mathbf{u}) e^{-(f(\mathbf{u}) - \mu_k)^2 / \sigma_k^2}$. In our settings, we modify the original MGS to a collaborative task. We

¹The MGS environment is at <https://github.com/Sonkyunghwan/QTRAN>

²The code is at <https://github.com/openai/multiagent-particle-envs>

³The code is at <https://github.com/google-research/football>

Algorithm 1: The optimization of GCS.

Ensure graph-based coordinated policy π^i and graph generator ρ ;
Initialize ϵ, γ, M, C , and the replay buffer $\mathcal{D} \leftarrow \emptyset$;
Initialize the parameters θ^i, ϕ^i for the graph-based coordinated policy networks, where $i = 1, \dots, n$, and φ for the graph generator network;
Initialize the target networks $\theta^{i-} = \theta^i$ and $\phi^{i-} = \phi^i$;
for each episode **do**
 Initial state $\leftarrow \{o_0^i\}_{i=1}^n$; // drop i as \mathbf{o}_0 for clarity;
 Initialize $h_0^{(1)}, \dots, h_0^{(n)}$ for RNN states;
 for each timestep t **do**
 Get $\mathcal{G}_t \sim \rho(\cdot | \mathbf{o}_t)$;
 Get the topological order $A_t = f(\mathcal{G}_t)$;
 for i in order A_t **do**
 augment the observations as
 $\hat{o}_t^i = (o_t^i, u_{t-1}^i, u_t^{pa(i)})$;
 sample the action with ϵ -greedy from
 $\mathcal{Q}_t^{\pi^i}(\hat{o}_t^i, u_t^i; \theta^i)$;
 Update RNN state $h_t^{(1)}, \dots, h_t^{(n)} \leftarrow h_{t-1}^{(1)}, \dots, h_{t-1}^{(n)}$;
 Receive reward \mathbf{r}_t and observe next state \mathbf{o}_{t+1} ;
 Add transition $\{\mathbf{o}_t, \mathbf{u}_t, \mathbf{r}_t, \mathbf{o}_{t+1}\}$ into \mathcal{D} ;
 end
 end
if episodes $> M$ **then**
 Sample a minibatch $\mathcal{B} = \{\mathbf{o}_j, \mathbf{u}_j, \mathbf{r}_j, \mathbf{o}_{j+1}\}_{j=0}^M \sim \mathcal{D}$;
 Update the policy π^i using \mathcal{B} and (4);
 Update the graph generator ρ using \mathcal{B} and $\nabla_\varphi L(\varphi)$;
 Update ϕ^i using (6);
 Every C steps reset $\theta_i^- = \theta_i$ for $i = 1, \dots, n$;
 Update the Lagrange penalty ξ and the multipliers λ_1, λ_2 ;
end
end

use two domains $[(\mu_1 = 5, \sigma_1 = 1.25), (\mu_2 = -5, \sigma_2 = 1.25)]$. The computation of the joint reward is as follows:

$$G(\mathbf{u}) = f(\mathbf{u}) e^{-(f(\mathbf{u}) - \mu_1)^2 / \sigma_1^2} - f(\mathbf{u}) e^{-(f(\mathbf{u}) - \mu_2)^2 / \sigma_2^2}. \quad (16)$$

Figure 3a shows the reward curves for this setting. According to the above definition, the reward is maximized when the resources of all the agents reach μ_1 or μ_2 . In this environment, the social welfare depends on the intensity of collaboration.

Cooperative Navigation (CN). Cooperative Navigation is a classic scenario implemented in the multi-agent particle world. This scenario has n agents and n landmarks, which are initialized with random locations at the beginning of each episode. The objective of the agents is to cooperate to cover all landmarks by controlling their velocities with directions. The action set \mathcal{A} includes five actions: [up, down, left, right, stop]. Each agent can only observe its velocity, position, and displacement from other agents and the landmarks. The shared reward is the negative sum of displacements between each landmark and its nearest agent. Agents must also

avoid collisions, since each agent is penalized with a ‘-1’ shared reward for every collision with other agents. We set the length of each episode as 25 time-steps. Therefore, the agents have to learn to navigate toward the landmarks cooperatively to cover all positions quickly and accurately. Figure 3b shows a classic scenario in Cooperative Navigation with $N = 3$.

Google Research Football (GRF). GRF is a realistically complicated and dynamic simulation environment without any clearly defined behavior abstraction, which is a suitable testbed for studying multi-agent decision making and coordination. In GRF, we use the *Floats* wrappers to represent the state. The *Floats* representation contains a 115-dimensional vector that summarizes the information, such as the ball position and possession, coordinates of all players, and the game state. Each player has 19 actions to control, including the standard move actions and different ball-kicking techniques. The rewards include the *SCORING* reward (-1, +1) and the *CHECKPOINT* reward, which is the shaped reward that specifically addresses the sparsity of *SCORING*. Detailed descriptions are shown in Appendix B.




Figure 3: The schematics of our experimental environments.

Baselines. We compared our results with several baselines as follows. VDN and QMIX are state-of-the-art value factorization approaches that follow the regime of centralized training and decentralized execution, belonging to the class of fully decentralized control policies, with which it is difficult to obtain coordinated behaviours. DCG uses fully connected graphs for belief propagation, which only allows the message passing of paired agents. DGN aims at learning abstract representations to make simultaneous decisions.

- VDN [7]: Value Decomposition Network (VDN) imposes the structural constraint of additivity in the factorization, which represents Q_{tot} as a sum of individual Q-values.
- QMIX [8]: This was proposed to overcome the limitation that VDN uses the linear decomposition and ignores any extra state information available during training. QMIX enforces Q_{tot} to be monotonic in the individual Q-values Q^i .
- DCG [28]: Deep Coordination Graph (DCG) factorizes the joint value function of all agents according to a coordination graph into payoffs between pairs of agents, which coordinates the actions between agents explicitly.
- DGN [30]: DGN relies on a graph convolutional network to model the relation representations, implicitly modeling the action coordination.

5.2 Main Results

Here we report the experimental results from the setup described in Section 5.1. Performance validation indicates the superiority of introducing ACG to multi-agent systems.

Collaborative Gaussian Squeeze. In this game, there are 10 agents, and the maximum episode length is also set to 10. To emphasize the feasibility and effectiveness of our proposed framework, we first conduct the experiment on CGS. We report the average episode rewards over 10 random runs, shown in Figure 4. Our proposed algorithm GCS outperforms the baseline methods by a large margin. It can be seen that our algorithm handles the collaborative problem well; the action coordination graph facilitates behavior learning to promote cooperation. Next, we will verify the effectiveness of our algorithm on more complicated environments.




Figure 4: Average episode rewards compared with baselines on Collaborative Gaussian Squeeze.

Cooperative Navigation. As shown in Figure 3b, Cooperative Navigation is a fully cooperative environment in which N agents (circles) must cooperate to reach N landmarks (crosses) with as few collisions as possible. We conduct experiments in Cooperative Navigation with $N = 4$ and with $N = 6$. Figures 5a and 5b show the learning curve comparisons for the two cases. We report the average episode reward at a training step, averaged over 10 independent running seeds.




Figure 5: Average episode rewards for four agents and six agents in Cooperative Navigation.

First, our algorithm outperforms most baseline algorithms by giving higher converged rewards. We consider that our performance improvement results from the action coordination graph representing the action dependency for better coordination. Moreover, our

algorithm converges fast during training, which is possibly because the hierarchical decision policies can efficiently induce coordination among agents in this cooperative setting. In addition, our algorithm can achieve a lower variance than those baselines, which indicates that the learned action coordination graph can reduce the uncertainty in decision making to facilitate cooperative behaviors among agents.

In contrast, VDN and QMIX take actions simultaneously without considering the action dependency among agents. They are faster during training, but that is of no benefit in inducing cooperation among agents. Additionally, DCG exhibits mediocre performance in this task. We believe that DCG considers only pairwise relationships between agents, which may disturb the overall balance in the system. In this case, DGN shows good performance consistent with ours, which shows that implicit action coordination modeling is also effective in pure cooperative settings.




Figure 6: Illustration of the effects of the action coordination graph. (a) Current state. (b) Topological structure of the learned ACG. (c) The intention of agents according to the ACG. The dashed lines represent the targets, and the solid lines represent the actions.

In order to clearly show the meaningful effect of the action coordination graph in the decision-making process, we give a visualization at a time step in an episode of Cooperative Navigation, as shown in Figure 6. Figure 6b shows the topological structure of the learned ACG, and it suggests that the decision dependency of the agent is [3, 4, 2, 1]. The parent sets of the four agents are denoted as $pa(1) = \{4, 2\}$, $pa(2) = \{3, 4\}$, $pa(3) = \emptyset$, and $pa(4) = \{3\}$, respectively. First, agent 3 decides to move to the bottom-left landmark, then agent 4 takes the best response and decides to move to the bottom-right in order to avoid conflict with agent 3. After agent 2 knows the decisions of the previous two agents, it chooses the closer upper-left as its target instead of the bottom-right. Finally, agent 1 moves after observing the decisions of agents 2 and 4. This visualization shows how agents’ joint actions deriving from the ACG representing the underlying decision dependencies achieves efficiency.

Google Research Football (GRF). To evaluate our method in complicated and dynamic environments, we conduct several experiments on GRF, as shown in Figure 7. In the 3-vs-2 scenario, three of our players try to score from the edge of the box, and the opponent team contains one defender and one keeper. In the 3-vs-6 scenario, there are six opponent players on the pitch to play against three of our players. In the 5-vs-5 scenario, each team has a keeper, an offensive player, and three defenders. Here, we report the average

episode reward at a training step for each scenario, averaged over 10 independent running seeds.

As can be seen in Figure 7, our algorithm always obtains higher rewards than all the baselines in the different scenarios of GRF. This indicates that our method is quite general in complicated and dynamic environments. Moreover, this performance improvement in GRF demonstrates that our approach is good at effectively handling stochasticity and sparse rewards. This is because the learned ACG with the decision dependency is an efficient way to mitigate uncertainty and induce cooperation among agents. Taking the 3-vs-6 scenario for further demonstration, the training curve of QMIX fluctuates and is unstable, indicating this method’s inability to adapt to the dynamically complicated scenario with multiple opponent players. Here, DCG shows a trend of non-convergence, but our algorithm steadily rises to converge and obtains the highest reward, which exhibits the modeling supremacy of our approach for handling complicated tasks.

5.3 Results on DAG Depth

We observe that the inference efficiency and the performance gains are inversely affected by the ACG’s depth. Therefore, we aim to find the suitable depth k that best balances the tradeoff. As shown in Figure 8, to validate the impact of the depth, we test our method on the Collaborative Gaussian Squeeze with different depth sizes of the learned ACG. In this figure, the horizontal axis is the depth, and the vertical axis is the testing episode reward averaged over five seeds. We test 1000 episodes for each seed and obtain an average episode reward.

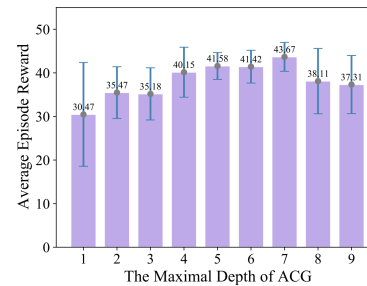


Figure 8: Effect of different DAG depth constraints on Collaborative Gaussian Squeeze.

As the depth of the ACG increases, the training time will increase correspondingly. However, the performance growth will gradually slow down, and performance degradation may even occur. In this case, $k = 5$ is the optimal depth that balances the computational burden and the performance gains. As for the reason for the performance degradation with $k = 8$ and $k = 9$, we speculate that as the hierarchy of action dependency deepens, the complexity of the hypothesis space for the inference will increase, and it becomes harder to learn the optimal policy. In summary, a higher dependency level of the graph structure can provide more decision information to promote coordination and facilitate performance. However, this higher dependency level leads to a lower efficiency of inference, as the leaf node on the ACG needs to wait for all the parent nodes’ decisions before it makes its own decision.




Figure 7: Average episode reward vs. training steps for comparisons with the baselines on Google Research Football.

5.4 Results on Dropping Edges

In this section, we verify the stability of the learned ACG in our proposed algorithm. Given a trained model with $depth = 5$ on Collaborative Gaussian Squeeze, we evaluate 1000 episodes and count the average number of edges of ACG, denoted as $edges_{num} \approx 28$. A fair comparison requires the same depth and number of edges during training and evaluation. Therefore, we generate a fixed DAG structure $\mathcal{G}_{\{5,28\}}$, see Appendix C, whose depth is 5 and whose number of edges is 28, as the baseline to compare with our algorithm.

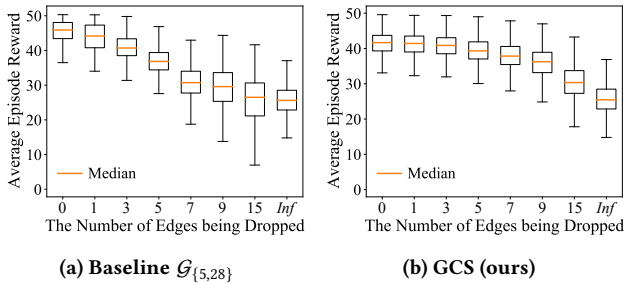


Figure 9: Box plots showing the distribution of the testing episode rewards in Collaborative Gaussian Squeeze.

In Figure 9, the horizontal axis represents the number of edges dropped from the generated or fixed graph structure, denoted as $\#drop$, and the vertical axis is the testing average episode reward over 1000 episodes. Inf denotes dropping all the edges. The box plots visually show the distribution of the testing episode rewards and skewness by displaying the data quartiles. From the overall trend of Figure 9, we can observe that the data quartiles of the baseline reduce faster and change more drastically than our algorithm when the disturbance of edge dropping increases. This demonstrates that our algorithm has better stability. Moreover, in Table 1, our algorithm outperforms the baseline with higher mean rewards in most cases, which demonstrates the power of the ACG learned by our model to promote the coordination among agents reliably and stably even when confronted with disturbances of different intensities. It is worth noting that to guarantee stability, a slight performance loss may occur. It will be an interesting future research direction to study the stability–performance trade-off.

Table 1: Mean value of testing episode rewards when dropping corresponding numbers of edges.

Methods	#drop							
	0	1	3	5	7	9	15	inf
The baseline	45.3	44.0	40.7	36.9	30.8	29.3	25.8	26.2
GCS (ours)	41.6	41.3	40.9	39.4	37.9	36.0	30.5	26.1

In summary, comparisons with VDN, QMIX, and DCG on three environments demonstrate that our algorithm achieves better performance, stronger stability, and more powerful modeling capability for handling dynamically complicated tasks than any of these methods. Moreover, the proposed DAG depth constraint provides an insightful view on balancing efficiency and performance.

6 CONCLUSIONS

In this paper, we introduce a novel graph generator and graph-based coordinated policy in MARL to dynamically represent the underlying decision dependency structure and facilitate behavior learning, respectively. We propose the DAGness-constrained and DAG depth-constrained optimization to balance training efficiency and performance gains. Extensive empirical experiments on Collaborative Gaussian Squeeze, Cooperative Navigation, and Google Research Football, as well as comparisons to baseline algorithms, demonstrate the superiority of our method.

Future research may consider improving the limited performance by upgrading the graph generator model. We will also investigate an automatic mechanism for finding an appropriate depth for the action coordination graph.

ACKNOWLEDGEMENTS

Jingqing Ruan is supported in part by the Strategic Priority Research Program of the Chinese Academy of Sciences under Grant XDA27010404 and in part by the National Nature Science Foundation of China under Grant 62073324. Co-author Haifeng Zhang is supported in part by the Strategic Priority Research Program of the Chinese Academy of Sciences, Grant No. XDA27030401.

REFERENCES

- [1] Oriol Vinyals, Timo Ewalds, Sergey Bartunov, Petko Georgiev, Alexander Sasha Vezhnevets, Michelle Yeo, Alireza Makhzani, Heinrich Küttler, John Agapiou, Julian Schrittwieser, et al. Starcraft ii: A new challenge for reinforcement learning. *arXiv preprint arXiv:1708.04782*, 2017.
- [2] Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. pages 6382–6393, 2017.
- [3] Lior Kuyper, Shimon Whiteson, Bram Bakker, and Nikos Vlassis. Multiagent reinforcement learning for urban traffic control using coordination graphs. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 656–671, 2008.
- [4] Joel Z Leibo, Vinicius Zambaldi, Marc Lanctot, Janusz Marecki, and Thore Graepel. Multi-agent reinforcement learning in sequential social dilemmas. In *International Conference on Autonomous Agents and Multiagent Systems*, pages 464–473, 2017.
- [5] Yali Du, Lei Han, Peng Sun, Jiechao Xiong, Qing Wang, Xinghai Sun, Han Liu, and Tong Zhang. Grid-wise control for multi-agent reinforcement learning in video game ai. In *International Conference on Machine Learning (ICML)*, pages 1–9, 2019.
- [6] Jiechuan Jiang and Zongqing Lu. Learning attentional communication for multi-agent cooperation. *Advances in Neural Information Processing Systems*, 31:7254–7264, 2018.
- [7] Peter Sunehag, Guy Lever, Audrunas Gruslys, Wojciech Marian Czarnecki, Vinicius Flores Zambaldi, Max Jaderberg, Marc Lanctot, Nicolas Sonnerat, Joel Z Leibo, Karl Tuyls, et al. Value-decomposition networks for cooperative multi-agent learning based on team reward. In *International Conference on Autonomous Agents and Multiagent Systems*, pages 2085–2087, 2018.
- [8] Tabish Rashid, Mikayel Samvelyan, Christian Schroeder, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. Qmix: Monotonic value function factorisation for deep multi-agent reinforcement learning. In *International Conference on Machine Learning*, pages 4295–4304, 2018.
- [9] Yali Du, Lei Han, Meng Fang, Tianhong Dai, Ji Liu, and Dacheng Tao. Liir: learning individual intrinsic reward in multi-agent reinforcement learning. In *Annual Conference on Neural Information Processing Systems*, pages 4403–4414, 2019.
- [10] Shariq Iqbal and Fei Sha. Actor-attention-critic for multi-agent reinforcement learning. In *International Conference on Machine Learning*, pages 2961–2970, 2019.
- [11] Haifeng Zhang, Weizhe Chen, Zeren Huang, Minne Li, Yaodong Yang, Weinan Zhang, and Jun Wang. Bi-level actor-critic for multi-agent coordination. In *AAAI Conference on Artificial Intelligence*, pages 7325–7332, 2020.
- [12] Dimitri Bertsekas. Multiagent rollout algorithms and reinforcement learning. *arXiv preprint arXiv:1910.00120*, 2019.
- [13] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 550:354–359, 2017.
- [14] Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, et al. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 588:604–609, 2020.
- [15] Duzhen Zhang, Tielin Zhang, Shuncheng Jia, Xiang Cheng, and Bo Xu. Population-coding and dynamic-neurons improved spiking actor network for reinforcement learning. *arXiv preprint arXiv:2106.07854*, 2021.
- [16] Xiaoqiang Wang, Yali Du, Shengyu Zhu, Liangjun Ke, Zhitang Chen, Jianye Hao, and Jun Wang. Ordering-based causal discovery with reinforcement learning. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 3566–3573. International Joint Conferences on Artificial Intelligence Organization, 8 2021.
- [17] Linghui Meng, Muning Wen, Yaodong Yang, Chenyang Le, Xiyun Li, Weinan Zhang, Ying Wen, Haifeng Zhang, Jun Wang, and Bo Xu. Offline pre-trained multi-agent decision transformer: One big sequence model tackles all smac tasks. *ArXiv*, abs/2112.02845, 2021.
- [18] Ming Tan. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *International Conference on Machine Learning*, pages 330–337, 1993.
- [19] Kyunghwan Son, Daewoo Kim, Wan Ju Kang, David Hostallero, and Yung Yi. QTRAN: learning to factorize with transformation for cooperative multi-agent reinforcement learning. In *International Conference on Machine Learning*, pages 5887–5896, 2019.
- [20] Jakob N. Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. Counterfactual multi-agent policy gradients. In *AAAI Conference on Artificial Intelligence*, pages 2974–2982, 2018.
- [21] Thomas Degris, Martha White, and Richard S Sutton. Off-policy actor-critic. *arXiv preprint arXiv:1205.4839*, 2012.
- [22] Lucian Buesoni, Robert Babuska, and Bart De Schutter. A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 38(2):156–172, 2008.
- [23] Jakob N Foerster, Yannis M Assael, Nando de Freitas, and Shimon Whiteson. Learning to communicate with deep multi-agent reinforcement learning. In *Annual Conference on Neural Information Processing Systems*, pages 2145–2153, 2016.
- [24] Sainbayar Sukhbaatar, Arthur Szlam, and Rob Fergus. Learning multiagent communication with backpropagation. In *Annual Conference on Neural Information Processing Systems*, pages 2252–2260, 2016.
- [25] Chongjie Zhang and Victor Lesser. Coordinating multi-agent reinforcement learning with limited communication. In *International Conference on Autonomous Agents and Multi-Agent Systems*, pages 1101–1108, 2013.
- [26] Peng Peng, Quan Yuan, Ying Wen, Yaodong Yang, Zhenkun Tang, Haitao Long, and Jun Wang. Multiagent bidirectionally-coordinated nets for learning to play starcraft combat games. *CoRR*, abs/1703.10069, 2017.
- [27] Yali Du, Bo Liu, Vincent Moens, Ziqi Liu, Zhicheng Ren, Jun Wang, Xu Chen, and Haifeng Zhang. Learning correlated communication topology in multi-agent reinforcement learning. In *International Conference on Autonomous Agents and Multi-Agent Systems*, pages 456–464, 2021.
- [28] Wendelin Böhrer, Vitaly Kurin, and Shimon Whiteson. Deep coordination graphs. In *International Conference on Machine Learning*, pages 980–991, 2020.
- [29] Sheng Li, Jayesh K Gupta, Peter Morales, Ross Allen, and Mykel J Kochenderfer. Deep implicit coordination graphs for multi-agent reinforcement learning. In *International Conference on Autonomous Agents and MultiAgent Systems*, pages 764–772, 2021.
- [30] Jiechuan Jiang, Chen Dun, Tiejun Huang, and Zongqing Lu. Graph convolutional reinforcement learning. *arXiv preprint arXiv:1810.09202*, 2018.
- [31] Xun Zheng, Bryon Aragam, Pradeep K Ravikumar, and Eric P Xing. Dags with no tears: Continuous optimization for structure learning. *Annual Conference on Neural Information Processing Systems*, pages 9472–9483, 2018.
- [32] Yue Yu, Jie Chen, Tian Gao, and Mo Yu. Dag-gnn: Dag structure learning with graph neural networks. In *International Conference on Machine Learning*, pages 7154–7163, 2019.
- [33] Yue Yu and Tian Gao. Dags with no curl: Efficient dag structure learning. In *Annual Conference on Neural Information Processing Systems*, 2020.
- [34] Sébastien Lachapelle, Philippe Brouillard, Tristan Deleu, and Simon Lacoste-Julien. Gradient-based neural dag learning. 2020.
- [35] Shengyu Zhu, Ignavier Ng, and Zhitang Chen. Causal discovery with reinforcement learning. In *International Conference on Learning Representations*, 2019.
- [36] Craig Boutilier. Planning, learning and coordination in multiagent decision processes. In *Theoretical Aspects of Rationality and Knowledge*, volume 96, pages 195–210, 1996.
- [37] I. N. Herstein. *Topics In Algebra*. John Wiley & Sons, 2nd edition, 1975.
- [38] Dimitri P Bertsekas. Nonlinear programming. *Journal of the Operational Research Society*, 48:334–334, 1997.
- [39] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *International Conference on Learning Representations*, 2018.
- [40] Chao Yu, Akash Velu, Eugene Vinitzky, Yu Wang, Alexandre Bayen, and Yi Wu. The surprising effectiveness of mapppo in cooperative, multi-agent games. *arXiv preprint arXiv:2103.01955*, 2021.
- [41] Karol Kurach, Anton Raichuk, Piotr Stańczyk, Michał Zajac, Olivier Bachem, Lasse Espeholt, Carlos Riquelme, Damien Vincent, Marcin Michalski, Olivier Bousquet, et al. Google research football: A novel reinforcement learning environment. In *AAAI Conference on Artificial Intelligence*, pages 4501–4510, 2020.