

Cooperative Policy Learning with Pre-trained Heterogeneous Observation Representations

Wenlei Shi
Microsoft Research
Beijing, China
Wenlei.Shi@microsoft.com

Xinran Wei
Beijing University of Posts and
Telecommunications
Beijing, China
weixinran@bupt.edu.cn

Jia Zhang
Microsoft Research
Beijing, China
Jia.Zhang@microsoft.com

Xiaoyuan Ni
Hong Kong University of Science and
Technology
Hong Kong, China
xniac@connect.ust.hk

Arthur Jiang
Microsoft ARD Incubation Team
Beijing, China
Shujia.Jiang@microsoft.com

Jiang Bian
Microsoft Research
Beijing, China
Jiang.Bian@microsoft.com

Tie-Yan Liu
Microsoft Research
Beijing, China
Tie-Yan.Liu@microsoft.com

ABSTRACT

Multi-agent reinforcement learning (MARL) has been increasingly explored to learn the cooperative policy towards maximizing a certain global reward. Many existing studies take advantage of graph neural networks (GNN) in MARL to propagate critical collaborative information over the interaction graph, built upon inter-connected agents. Nevertheless, the vanilla GNN approach yields substantial defects in dealing with complex real-world scenarios since the generic message passing mechanism is ineffective between heterogeneous vertices and, moreover, simple message aggregation functions are incapable of accurately modeling the combinational interactions from multiple neighbors. While adopting complex GNN models with more informative message passing and aggregation mechanisms can obviously benefit heterogeneous vertex representations and cooperative policy learning, it could, on the other hand, increase the training difficulty of MARL and demand more intense and direct reward signals compared to the original global reward. To address these challenges, we propose a new cooperative learning framework with pre-trained heterogeneous observation representations. Particularly, we employ an encoder-decoder based graph attention to learn the intricate interactions and heterogeneous representations that can be more easily leveraged by MARL. Moreover, we design a pre-training with local actor-critic algorithm to ease the difficulty in cooperative policy learning. Extensive experiments over real-world scenarios demonstrate that our new approach can significantly outperform existing MARL baselines as well as operational research solutions that are widely-used in industry.

KEYWORDS

Multi-agent Reinforcement Learning; cooperative policy learning; graph attention network for state representation; pre-training method

ACM Reference Format:

Wenlei Shi, Xinran Wei, Jia Zhang, Xiaoyuan Ni, Arthur Jiang, Jiang Bian, and Tie-Yan Liu. 2021. Cooperative Policy Learning with Pre-trained Heterogeneous Observation Representations. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Online, May 3-7, 2021*, IFAAMAS, 9 pages.

1 INTRODUCTION

Cooperative multi-agent systems (MAS) [22] have been employed to deal with a wide range of real-world applications, such as large scale fleet management [16], empty container repositioning [13], and traffic signal control [32]. Instead of following classical single-agent planning methods, cooperative MAS stands essentially as a more efficient paradigm to resolve the global optimization task through decentralized agents with a relatively-simplified agent-wise problem-solving scheme. Essentially, the success of such cooperative MAS highly relies on its capability to model the cooperation between agents with the purpose of pursuing a maximized global utility. In reality, to bolster cooperation among complex relationships between agents and environmental entities, it is indispensable for each agent to fully perceive the state information about other inter-correlated agents and environmental entities.

To this end, previous studies [11, 18] proposed to model the entire environment with an *interaction graph* where vertices represent the agents and environmental entities and edges exist when two vertices can interact with each other. And, they leveraged Graph Neural Networks (GNN) [7, 12, 30] to propagate critical cooperation signals between inter-connected agents within the Multi-Agent Reinforcement Learning (MARL) framework [8]. However, while the vanilla GNN model can perform well in relatively simple games, two major weaknesses limit its representability to capture the state




Figure 1: The overall structure of EncGAT-PreLAC. From left to right, it uses EncGAT for interaction representation and feeds the representations to the actor and critic headers. The overall model is first pre-trained with the local loss L^{LOC} and then fine-tuned with a global loss L^{LOC} which are defined in Formula (3) and Formula (4).

information through the interaction graph in many complex real-world scenarios.

First, the generic procedure of message propagation in vanilla GNN assumes that the vertices in the interaction graph are homogeneous. In practice, however, the vertices can be heterogeneous, which impedes the efficient message passing due to the incompatible feature spaces between heterogeneous vertices. Furthermore, another weakness lies in simply using pooling function, such as average and maximization, as message aggregation, which may cause significant information loss or the over-smoothing problem. Although more expressive pooling functions such as LSTM have been employed to enhance GNN structures [7], none of them, to our best knowledge, has been applied in MARL scenarios.

Henceforth, it becomes necessary to design a more expressive GNN with more informative message passing and aggregation mechanisms to increase the individual agent’s ability to grasp the complex representations of neighbors’ messages. Such enhancement over GNN, on the other hand, is likely to encounter a dilemma. Specifically, it may lead to drastically-soaring training difficulties, in terms of requiring either stronger training signals or much more experience, for each agent to converge into an effective cooperative policy. However, the signals measuring the accomplishment of the global optimization goals, e.g. the total value of served orders in a ride-sharing platform [16] or total shortage of empty containers in ocean transportation [13], are influenced by the entangled actions from multiple agents. One agent could hardly evaluate its own actions with limited experiences especially when other agents’ policy is also non-stationary. Merely leveraging such global signals and letting agents learn *from scratch* will lead to poor sample efficiency and make the training process unbearably slow.

To obtain effective information representation from the heterogeneous interaction graph and learn the cooperative policy for decentralized agents, in this paper, we propose a new multi-agent reinforcement learning framework as Figure 3. Particularly, this framework leverages the graph attention with a delicately designed encoder-decoder structure for efficient message passing and aggregation. Moreover, we introduce a pre-training technology to alleviate the difficulty of learning the cooperative policy with global rewards.

More concretely, to magnify the GNN’s representability in the process of information passing and aggregation, we employ an encoder-decoder based graph attention network (EncGAT). Similar to its previous success in the NLP domain, this approach can consecutively embed the sequence of neighboring vertices into a

high-level representation space with a self-attention encoder and then generate the aggregated feature through an attention decoder with the feature of the vertex itself as the query, whose details are demonstrated in Figure 2.

To ease the difficulty in learning cooperative policy for decentralized agents with global rewards, we propose a pre-training with a local actor-critic (PreLAC) algorithm. Specifically, a pre-training phase is designed to learn a selfish policy for each agent by using local rewards. Since local rewards, which are naturally existing in most global optimization scenarios (e.g. regional revenue of ride-sharing [16] and port-level container shortage [13]), are mainly affected by the corresponding individual agent and its neighbors, it is much easier to train a selfish policy, specified by a local actor and a local critic with the shared EncGAT layer, for each agent. Following this pre-training, a formal training phase will be carried out, in which decentralized agents with cooperative local actors and a shared global critic will be continuously trained, starting from using the pre-trained EncGAT layer, towards maximizing the global rewards. To avoid the training collapse, local critics are still working in the formal training phase for the parameters updating in EncGAT. Compared to the straightforward approach with a simple combination of global and local rewards, the PreLAC algorithm does not need to manually trade off between convergence and optimality, which can thus result in a more effective cooperative policy with less manual intervention.

To evaluate the effectiveness of our framework, we conduct experiments over a real-world problem, i.e., empty container repositioning, in the scenario of ocean transportation. Extensive experiments have shown that our new framework can significantly outperform either the operations research solution, a widely used industry standard, or existing GNN-powered MARL baselines. Further investigations have also been conducted to demonstrate the advantages of the proposed encoder-decoder attention structure and the introduction of the pre-training task in terms of their effectiveness in accelerating the convergence and improving the ultimate global performance.

Our main contribution can be summarized as follows:

- We propose an encoder-decoder based graph attention (EncGAT) layer for representation learning of the complicated interactions between heterogeneous entities in the MAS with global optimization goal.
- We introduce the pre-training with local actor-critic (PreLAC) algorithm for cooperative policy learning, in which the selfish policy learning with the purpose of local rewards maximization is taken as a pre-trained task and the pre-trained

EncGAT parameters will be used for further cooperative policy learning.

- Extensive experiments are conducted on the practical empty container repositioning (ECR) problem and the results demonstrate the advantages of the proposed framework.

In addition to the ECR problem, the proposed EncGAT-PreLAC framework in fact can be applied to a series of global optimization scenarios, where local rewards can be defined. These scenarios include but are not limited to vehicle re-balancing in ride-sharing, traffic control, and even football match.

2 RELATED WORK

Recent years have witnessed the increasing research attention on applying cooperative multi-agent systems into many application scenarios. Previously, traditional operational research (OR) methods were widely used to achieve global optimization in a centralized way [3, 24, 25]. However, the prerequisite of both the forecasting module and artificial constructions of mathematical objectives and constraints gave rise to unsatisfactory performance in practice since OR methods are difficult to be adapted to a dynamic environment. With the recent rapid development of reinforcement learning, a rising number of efforts have turned to formulate MAS into stochastic games and then solve them using the multi-agent reinforcement learning (MARL) approach. There have been a couple of corresponding practices, such as empty container repositioning [13], shared bike repositioning [6, 15, 21], ride-sharing fleet management [4, 16], traffic light control [32], and cloud resource allocation [17, 33].

Yet, in many real-world MAS with complex networks, one critical challenge, particularly in terms of how to feed each agent with the informative representation of a comprehensive observation scope for its decision-maker, has not been fully addressed. Some previous studies simply produced such representation using pre-defined contextual statistics, which has to be coupled with specific scenarios and apparently limits the generality. Most recently, GNN has been increasingly used to learn and enhance the representation of observation scope in the non-Euclidean space. For instance, MAGnet [18] learns relevance information from the observation in the form of a relevance graph, where relation weights are learned by pre-defined loss function based on heuristic rules. The relation weights in DGN [11] are learned with a temporal regularization under the assumption that the attention of one agent does not change much in two adjacent time steps, which is not always true in practical problems yet. Another work [1] leverages the inductiveness of graph attention to learn a transferable policy. There have also been recent efforts to use GNN to enhance collaboration in a practical problem called traffic signal control [2, 32].

However, all the studies above treat the observation as an isomorphic graph, ignoring the heterogeneity of agents and environmental entities commonly existed in real-world scenarios. A most recent work [26] exploits the sparsity of interactions between agents and tackles the problem of heterogeneity of agents from different teams. Other than [26], in this paper, we focus on the gap of feature space between both agents and environmental entities and design a new graph attention based aggregation approach.

There are also some existing studies paying attention to enhancing the training signal. Counterfactual advantage [5] is used for

credit assignment to filtering out contributions of other agents. Cooperative reward mechanism design [13] is used to compute the cooperative local reward. Value decomposition [23, 28] is also a hot topic to solve the problem of cooperative MARL with a single joint reward signal. In this paper, we leverage the direct local reward to accelerate the training, but their methods can be integrated into our framework.

In the field of GNN, our work can be viewed as a graph-level training task with vertex-level pre-training. Recently, pre-training GNNs in supervised or semi-supervised tasks have also been studied in several works[9, 20], but our pre-training pipeline is quite different from theirs. In addition, there are several studies[31] about the heterogeneous graph. However, our proposed networks introduce a different encoder-decoder structured aggregation function. Besides, our work aims to solve the sequential decision-making task, which basically yields a different problem domain other than theirs.

3 PROBLEM FORMULATION

Heterogeneous interaction graph. In many practical MAS, agents need to cooperate to achieve a maximized global utility. An interaction graph can be constructed to characterize the interaction between agents, e.g. fleets [16] and ports [13], and environment entities, e.g. wooden walls [18]. Specifically, the vertices, in the interaction graph, are the agents and environmental entities, while edges stand for their corresponding interactions. Because agents and environmental entities usually yield different types, the interaction graph is essentially a heterogeneous graph. We follow the previous work [27] to formally define the heterogeneous graph.

Definition 3.1. A heterogeneous interaction graph, denoted as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, is a graph including a set of vertices \mathcal{V} and a set of directed edges \mathcal{E} , where $\mathcal{V}^* \subset \mathcal{V}$ is the set of agents. In \mathcal{G} , we denote the set of the types of vertices and edges as χ and ψ , and denote the type of vertex i and edge j as χ_i and ψ_j . The set of the types of edges that are connected to vertex v_i are denoted as $\mathcal{N}_{\psi} : v_i \rightarrow \mathbb{D}$. Then we define the mapping function from the vertex to the set of neighbors that connected by the edges of type $d \in \mathcal{N}_{\psi}(v_i)$ as $\mathcal{N}_d : v_i \rightarrow \mathbb{V}$.

MAS with heterogeneous interaction graph. We formally model the MAS with heterogeneous interaction graph as a Semi-POMDP $\langle \mathcal{V}^*, \mathcal{A}, \mathcal{S}, \mathcal{O}, \mathcal{R}, \mathcal{P}, \gamma \rangle$. More concretely, \mathcal{V}^* denotes the aforementioned agent set; \mathcal{A} denotes the action set specified by the joint action space of all agents, i.e., $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2 \times \dots \times \mathcal{A}_{|\mathcal{V}^*|}$; \mathcal{S} represents the global state that can be defined by a heterogeneous interaction graph \mathcal{G} comprised of the agent vertices, environmental entity vertices and interaction edges; \mathcal{O} represents the set of each agent's observations, which is the joint observation space, i.e., $\mathcal{O} = \mathcal{O}_1 \times \mathcal{O}_2 \times \dots \times \mathcal{O}_{|\mathcal{V}^*|}$, where \mathcal{O}_i is the observation space of agent i and contains the local vertices and interaction edges near the agent; \mathcal{R} is the global reward function, which is defined as $\mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$; note that, in most of MAS, there exists local rewards for each agent¹, and we can similarly define the local reward function $R_i^{\text{LOC}} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$; \mathcal{P} is the transition probability function $\mathcal{S}_t \times \mathcal{A} \rightarrow \Delta(\mathcal{S}'_{t+k})$, where Δ is the probability simplex over the

¹For examples, scores of a football player, orders completed by a driver and shortage of containers at a port.

next global state and the time interval k is non-constant time interval from current state to the next state; γ is the discount factor on the unit time interval. The policy of agent i is defined as a function $\pi_i : O_i \rightarrow \mathcal{A}_i$. The joint policy is defined as $\pi = (\pi_1, \dots, \pi_{|\mathcal{V}^*|})$.

4 PROPOSED FRAMEWORK

We introduce our framework based on the multi-agent A2C algorithm for convenience. In fact, our framework can be easily extended to other actor-critic based algorithms. As is shown in Figure 1, the proposed EncGAT-PreLAC framework is composed of two parts: an encoder-decoder graph attention (EncGAT) model and a pre-training procedure with local actor-critic (PreLAC). The EncGAT model is designed to learn the informative representation of the interaction graph, which will be fed into both the actor and critic networks to facilitate learning the cooperative policy. However, introducing the delicate-designed EncGAT model into the A2C algorithm makes the learning process more difficult in MAS. To improve both the learning efficiency and the effectiveness of the policy, we first leverage the local reward to train multiple selfish local actors and critics with the EncGAT model. Then, we initialize the final multiple actors and the global critic with the pre-trained EncGAT and conduct the fine-tuning with the global reward. More details are introduced in the following sub-sections.

4.1 Encoder-Decoder Based Graph Attention for State Representation

As aforementioned, the vertices heterogeneity together with the intricate interactions between them prevents vanilla GNN to effectively learn the informative representations from the complex interaction graph. The EncGAT model is designed to solve these problems.




Figure 2: Encoder-decoder to aggregate interaction information from neighbors which is the basic block in EncGAT.

The encoder-decoder based attention procedure is shown in the figure 2. For vertex (agent) $v_i \in \mathcal{V}^*$ with feature vector x_i , its neighbor set with type d is denoted as $\mathcal{N}_d(v_i)$ and the corresponding feature matrix is \mathbf{x}_i^d . The calculation procedure of the encoder is based on self-attention, which is similar to the transformer block and shown below:

$$\begin{aligned} \mathbf{a}_i^d &= \text{ATT} \left(W_{ec}^Q \mathbf{x}_i^d, W_{ec}^K \mathbf{x}_i^d, W_{ec}^V \mathbf{x}_i^d \right) \\ \mathbf{k}_i^d &= \text{LN} \left(\mathbf{a}_i^d + \mathbf{x}_i^d \right) \\ \mathbf{z}_i^d &= \text{RELU} \left(\mathbf{k}_i^d W_1^T + b_1 \right) W_2^T + b_2 \\ \mathbf{o}_i^d &= \text{LN} \left(\mathbf{k}_i^d + \mathbf{z}_i^d \right) \end{aligned} \quad (1)$$

where ATT is the scaled dot product attention which is the same as that in Transformer [29], W_{ec}^Q, W_{ec}^K and W_{ec}^V are the projection

matrices of query, key and value in the attention function, W_1 and W_2 are the square matrices of the linear layers, b_1 and b_2 are the bias parameters, LN is the layer norm and RELU is the ReLU activation function. The final embedding \mathbf{o}_i^d is already transformed into the same vector space with x_i in case that v_i is not of type d . We the whole encoder process as the function ENC. Through the function ENC, the feature of each neighboring vertex of v_i is represented by considering other neighbors' information in the ATT function.

The aggregation among the same type of neighbors is conducted through the decoder attention as is shown in the left part of the figure 2. In decoder attention, another scaled dot product attention is applied, which uses the feature vector of v_i , i.e. x_i , as the query, and the encoded neighbors' features matrix \mathbf{o}_i^d as the keys and values. To align the dimension, the feature vector x_i is first expanded as a matrix \mathbf{x}_i . The mathematical description is shown as follows:

$$\begin{aligned} \mathbf{e}_i^d &= \text{ATT} \left(W_{dc}^Q \mathbf{x}_i, W_{dc}^K \mathbf{o}_i^d, W_{dc}^V \mathbf{o}_i^d \right) \\ \mathbf{f}_i^d &= \text{LN} \left(\mathbf{e}_i^d + \mathbf{o}_i^d \right) \\ \mathbf{c}_i^d &= \text{RELU} \left(\mathbf{f}_i^d W_3^T + b_3 \right) W_4^T + b_4 \\ \mathbf{h}_i^d &= \text{LN} \left(\mathbf{f}_i^d + \mathbf{c}_i^d \right) \end{aligned} \quad (2)$$

where W_{dc}^Q, W_{dc}^K and W_{dc}^V are the projection matrices of query, key and value in the attention function, W_3 and W_4 are the square matrices of linear layers and b_3 and b_4 are the bias parameters. Then we squeeze the expanded dimension of \mathbf{h}_i into the vector h_i

Finally, as the neighbors of vertex v_i may have different types, we simply concatenate all the output h_i^d as the final feature of vertex v_i . The expression is as $h_i = \left\| \left\{ h_i^d \mid d \in \mathcal{N}_\psi(v_i) \right\} \right\|$.

4.1.1 Discussion on EncGAT. Indeed, we provide a new perspective on the cooperative information aggregation within a heterogeneous graph through EncGAT. To cooperate with other agents, an agent must perceive the heterogeneous information from neighbors. The success of applying the self-attention architecture in the natural language processing tasks inspires us to leverage the same idea to enforce cooperative information aggregation. For one thing, it enables the agent to parse the contextualized interactions with each individual neighbor given all other neighbors. Moreover, the encoder-decoder attention can eliminate the feature gap between two heterogeneous vertices by mapping them into a unified representation space and reduce the difficulty in cooperative information extraction. The vanilla graph attention in fact can be seen as a simplified case without the encoder, which considers the interaction with each neighbor independently and then aggregates with a simple summation. This can only model the simplest linear combinational interactions of neighbors and will fail to model complex combinational interactions such as the relative order of consumer and producer in the inventory management, as is stated in Section 1. The self-attention we use in the encoder can be seen as weaving a complete graph between the neighbors of the same type and conducting the graph attention in these 'sub-graphs'. With the aid of self-attention in the encoder, the feature of each neighbor is calculated based on contextualized information from other neighbors such that it can capture the complex reactions in the interactions.

In addition, the encoder-decoder structure we state here differs from the Transformer [29] in that the decoder in our model only gets a single vector as query, rather than a sequence. Therefore, no self-attention on the query is used in the decoder. The recent work GTR [14] uses transformer blocks in graph attention, which directly adds up features of the vertex and the weighted neighborhood in inter-graph message passing. It can be seen as the decoder-only version of EncGAT.

4.2 Pre-Training with Local Actor-Critic

A crucial challenge of training the EncGAT lies in that a direct optimization towards the global reward may lead to insufficient observation understanding, which will slow down the convergence and even erode the performance of the learned cooperative policy. In this paper, we manage to address this problem by proposing the pre-training with the local actor-critic (PreLAC) algorithm.

In the local actor-critic task, each agent is trained to maximize its own interest, that is the discounted sum of its local reward $\{r_{i,t}^{\text{LOC}} \mid 1 \leq i \leq |\mathcal{V}^*|\}$, where t is the time-step. It encourages each agent to learn the concrete business logic about how the local reward is generated and how the neighborhood interactions affect the reward generation.

As is shown in figure 1, the EncGAT network is used as a shared embedding module for the actor and critic networks in the pre-training task. It receives the vertex-level training signals from the local actor and local critic headers. The Formula (3) shows the training loss in this stage.

$$\begin{aligned} \hat{A}_{i,t}^{\text{LOC}} &= r_{i,t+1}^{\text{LOC}} + \dots + \gamma^{k-1} r_{i,t+k}^{\text{LOC}} + \gamma^k V^{\pi_i}(O_{i,t+k}) - V^{\pi_i}(O_{i,t}), \\ \delta_{i,t}^{\text{LOC}} &= \left(\hat{A}_{i,t}^{\text{LOC}}\right)^2, \phi_{i,t}^{\text{LOC}} = -\log \pi_i(O_{i,t}, A_{i,t}) \hat{A}_{i,t}^{\text{LOC}}, \\ L^{\text{LOC}} &= \lambda_l^{\text{LOC}} \sum_i \sum_t \delta_{i,t}^{\text{LOC}} + \sum_i \sum_t \phi_{i,t}^{\text{LOC}}, \end{aligned} \quad (3)$$

where $r_{i,t}^{\text{LOC}}$ denotes the local reward of agent v_i at time step t , while k refers to the interval from current state to the next state, following the Semi-POMDP setting in Section 3; $V^{\pi_i}(\cdot)$ is the local value function under the local policy π_i of v_i ; $A_{i,t}^{\text{LOC}}$ is the action of agent v_i ; and the hyper-parameter λ_l^{LOC} is the weight of the critic loss. As we can see, the policy loss $\phi_{i,t}^{\text{LOC}}$ and the critic loss $\delta_{i,t}^{\text{LOC}}$ are computed based on the advantage function $\hat{A}_{i,t}^{\text{LOC}}$ derived from the local critic. The training follows the vanilla Actor-Critic algorithm, thus the $\hat{A}_{i,t}^{\text{LOC}}$ in the third formula are detached in back-propagation and hence will not affect the update of the critic headers.

In the fine-tuning stage, we reinitialize the local actor headers for each agent and create a new global critic header, while the parameters of EncGAT and local critic headers are inherited from the pre-trained part. Different from the pre-training, the global reward r_t is used as a training signal for the local actors' network (including the shared EncGAT and actor header of each agent) and the global critic network, which encourages cooperation towards the global goal. Meanwhile, to prevent the training collapse and avoid the interference of random initialed parameters, we keep using the local critic loss in this stage to supervising the update of the shared EncGAT part with the agent-level local loss derived from local critics. Therefore, the loss function can be formally defined as

the linear combination of three parts: the local critic loss δ_t^{LOC} , the global critic loss δ_t and the actor loss ϕ_t . That is

$$\begin{aligned} \hat{A}_t &= r_{t+1} + \dots + \gamma^{k-1} r_{t+k} + \gamma^k V^\pi(S_{t+k}) - V^\pi(S_t), \\ \delta_t^{\text{LOC}} &= \sum_i \delta_{i,t}^{\text{LOC}} = \sum_i \left(\hat{A}_{i,t}^{\text{LOC}}\right)^2, \delta_t = \hat{A}_t^2, \\ \phi_t &= -\left(\sum_i \log \pi_i(O_{i,t}, A_{i,t})\right) \hat{A}_t, \\ L &= \lambda_l \sum_t \delta_t^{\text{LOC}} + \lambda_g \sum_t \delta_t + \sum_t \phi_t, \end{aligned} \quad (4)$$

where r_t denotes the global reward, \hat{A}_t represents the global state advantage, V^π is the global value function under the joint policy π and the hyper-parameter λ_l and λ_g are the weights of the local critic loss and the global critic loss. Similar to the pre-training stage, the gradient of the advantage \hat{A}_t is detached in the training and will not affect the update of the global critic header.

4.3 Algorithm of EncGAT-PreLAC Framework

The whole EncGAT-PreLAC framework can be summarized as Algorithm 1. For simplicity of description, we assume all agents belong to the same type while other vertices in the interaction graph can be heterogeneous in the algorithm. Meanwhile, the actor header and critic header are shared among all agents, which can be overwritten according to actual demand. In addition, our framework supports both the synchronous and asynchronous decision of agents, thus we denote \mathcal{V}_t^* as the set of actionable agents at the time step t in both the pre-training and the training procedure.

5 EXPERIMENTS

We evaluate our framework on the empty container repositioning problem (ECR) in ocean transportation, which cost billions of dollars per year for shipping companies. All the simulators and experiments in our work are based on the open-sourced platform Multi-Agent Resource Optimization (MARO) [10].

5.1 Task Description

Empty containers are the core resource used to load goods in ocean transportation. The serious imbalance of import and export among countries and regions all around the world results in extreme imbalance of supply and demand (SnD) for empty containers. With repositioning operation, some ports will stack large amount of containers with high storage cost while some other ports have no containers to satisfy customers' shipping requirements. Obviously, both cases will cause serious losses for the shipping companies and hinder the global trade. Empty container repositioning is to leverage the remaining capacity in vessels to explicitly transfer empty containers between deficit ports and surplus ports. Each vessel sails around a route and the travel time is influenced by unpredictable factors such as weather and ocean current. When a vessel arrives at a port, it first loads and discharges cargoes, and then triggers the repositioning action of the port to load to or discharge from the vessel a certain amount of empty containers.

We model the ECR problem as a MAS as follows. We assign an agent to each port to manage the containers. When a vessel arrives

Algorithm 1 EncGAT-PreLAC Framework

```

1: procedure PRE-TRAINING
2:   Initialize parameters of EncGAT, actor headers and critic
   headers:  $W_e^{\text{PRE}}, W_a^{\text{PRE}}$  and  $W_c^{\text{PRE}}$ 
3:   Initialize the experience pool  $D^{\text{PRE}} = \emptyset$ 
4:   Initialize simulator env
5:   for  $k$  in  $1, \dots, \text{MAX}$  do
6:     while not env.done() do
7:        $\{(O_{i,t}, \mathcal{A}_{i,t}) \mid i \in \mathcal{V}_t^*\} \leftarrow \text{env.get\_state}()$ 
8:       for  $v_i \in \mathcal{V}_t^*$  do
9:          $A_{i,t} \sim \pi_i(A \in \mathcal{A}_{i,t} \mid O_{i,t}; W_e^{\text{PRE}}, W_a^{\text{PRE}})$ 
10:         $(r_t, \{r_{j,t}^{\text{LOC}} \mid j \in \mathcal{V}^*\}) = \text{env.step}(\{A_{i,t} \mid i \in \mathcal{V}_t^*\})$ 
11:        Store experience to  $D^{\text{PRE}}$ 
12:      for batch in  $D^{\text{PRE}}$  do
13:        Compute the loss  $L^{\text{LOC}}$  according to formula (3)
14:        Update  $W_e^{\text{PRE}}, W_a^{\text{PRE}}, W_c^{\text{PRE}}$  with the gradient of  $L^{\text{LOC}}$ 
15:        Reset the simulator env.reset()
16:      Set  $D^{\text{PRE}} = \emptyset$ 
17: procedure TRAINING
18:   Initialize the parameters of EncGAT  $W_e$  with  $W_e^{\text{PRE}}$ 
19:   Initialize the parameters of the local critic header  $W_c^{\text{LOC}}$ 
   with  $W_c^{\text{PRE}}$ 
20:   Initialize the parameters of the actor header  $W_a$  and the
   global critic header  $W_c$  with random value
21:   Initialize the experience pool  $D = \emptyset$ 
22:   Initialize simulator env
23:   for  $k$  in  $1, \dots, \text{MAX}$  do
24:     while not env.done() do
25:        $\{(O_{i,t}, \mathcal{A}_{i,t}) \mid i \in \mathcal{V}_t^*\} \leftarrow \text{env.get\_state}()$ 
26:       for  $v_i \in \mathcal{V}_t^*$  do
27:          $A_{i,t} \sim \pi_i(A \in \mathcal{A}_{i,t} \mid O_{i,t}; W_e, W_a)$ 
28:          $(r_t, \{r_{j,t}^{\text{LOC}} \mid j \in \mathcal{V}^*\}) = \text{env.step}(\{A_{i,t} \mid i \in \mathcal{V}_t^*\})$ 
29:         Store experience to  $D$ 
30:       for batch in  $D$  do
31:         Compute the loss  $L$  according to formula (4)
32:         Update  $W_e, W_a, W_c^{\text{LOC}}, W_c$ 
33:         env.reset()
34:       Set  $D = \emptyset$ 

```

at the port, the corresponding agent is triggered to make a repositioning decision, which is in fact an event-driven reinforcement learning [13, 19]. The repositioning *action* space is $\{0, 1, \dots, 21\}$ where action 0 means discharging all empty containers on the vessel and action 21 means loading all empty containers of the port on board. Other actions are evenly distributed with a step of 10 percent. The *observation* of an agent is its own information and that of related ports and vessels. The optimization goal is to minimize the total shortage of empty containers to fulfill the transportation needs because of the SnD imbalance among ports, which is, on the contrary, to maximize the amount of the fulfilled demand. Therefore, we define the *global reward* as a homogeneous linear function of the number of empty containers used to fulfill transportation orders at a certain time period, that is $r_t = k \cdot \sum_i^n a_{i,t}$, where k is a positive scale factor, n is the number of ports and $a_{i,t}$ is the quantity

of consumed containers in port v_i at the time period t . At the same time, the most straightforward definition of the *local reward* for the port v_i is the value $k \cdot a_{i,t}$.

To realize the efficient repositioning, the agents should have the ability to extract useful information from the observation on the heterogeneous interaction graph. As the interactions happen when a vessel arrives at a port, the arrival time of the vessels from different routes plays an important role in decision making. For example, suppose there are two routes, the first one is in lack of containers and the second is rich in containers, namely the demand route and supply route. Two vessels from these routes will arrive at a port in particular time steps. Obviously, whether the vessel in the demand route arrives first or second will substantially affect the optimal repositioning decision of this port, e.g. reserving containers for the incoming vessel or using them out. Such a relative order of two vessels actually poses the requirements for our model to learn the complicate interactions among neighbors, which the vanilla pooling functions like averaging and maximizing will fail to capture as we state in Section 1. Another necessary feature for the agents is the cooperation between agents based on agents' comprehensive understanding of the observations. A common phenomenon is that the demand routes and supply routes are not necessarily connected and agents should learn to play as brokers to conduct the multi-hop repositioning.

5.2 Data and Simulator

To build the simulator for the ECR problem, we employ the public routes of an international transportation company, including 22 ports located in main trading countries, 13 routes connecting these ports, and 46 vessels sailing on the routes. The concrete settings of the route topology are attached in the supplementary material. The shipment data between any two ports is generated based on the inter-country trading statistics released by WTO in 2019. Because only the yearly statistics are available while the temporal resolution of the simulation is daily, we distribute the total transportation amount into each time-step by two alternative rules. The first is simply the constant called 'normal' setting and the second is to distribute the amount by a trend function called 'hard' setting. The transportation time between two ports is computed by their sea distance. Gaussian noises are added in both transportation time and trading amount to include the environmental uncertainty in each episode.




Figure 3: The network structure in ECR task.

5.3 Input and EncGAT implementation

In the ECR problem, there are two kinds of vertices in the interaction graph, ports and vessels denoted as $\{v_i^p\}$ and $\{v_j^s\}$ respectively. We establish edges from one port to other ports and vessels in the same route, which represents the transportation interactions. For the port, the input features include the quantities of the empty containers, the reserved empty containers, the laden, the laden on board, transportation orders, the capacity, and the remaining space. The statistical information is also included such as the quantities of the total orders, the fulfilled orders, the failed orders, and their accumulated values in the past 5 time-steps. The input features of the vessel are 7 variables which are the quantities of the empty containers, the laden, the capacity, the remaining space, the accumulated quantities of empty containers, laden, and their difference. As each agent has to learn to predict trend of the inventory and plan for the future, it can hardly determine the strategy with only the input of the current information. Therefore, we let each agent look back a period. We denote the input features of each vertex over the time period $[t - n : t]$ as $x_{i,[t-n:t]}$.

The network structure is shown in Figure 3. It follows the general design of the proposed framework, but we add three features based on the characteristics of the ECR problem. The first one is the temporal attention, which is used to capture the temporal information of the temporally sequential input, is added before the EncGAT module. For example for vertex v_i , the temporal embedding is computed as follows:

$$x_i = \text{ATT} \left(W_s^Q x_{i,t}, W_s^K x_{i,[t-n:t]}, W_s^V x_{i,[t-n:t]} \right) \quad (5)$$

where W_s^Q , W_s^K and W_s^V are the projection matrices for the query, key and value respectively. In the experiment, the lookback period n is set to 20. The second one is that to consider the edge feature like the arriving time and the distance, we concatenate it with the features of the corresponding neighbor before feeding it to EncGAT. The third one is that the representations of both the port and the current vessel are fed into the actor header since the vessel situation is important for decision making.

Finally, we stack two encoder-decoder attentions in the EncGAT and residual connections are used to prevent over-smoothing. The actor and critic headers consist of two fully-connected layers also with residual connections. We share the headers of actor and local critic between agents, which makes the overall framework inductive [7]. Empirical studies show the generality of policy with the parameter sharing, which is illustrated in the supplemental file.

5.4 Comparison with Baselines

In this section, we compare our framework with existing baseline methods with the total fulfillment ratio, that is the ratio of fulfilled demands to all demands, as the evaluation metric. To keep fairness, we search the learning rate for all deep learning algorithms and fix other hyper-parameters. We run 5 times for each one with the best learning rate. The results are shown in Table 1. It should be pointed out that since most of the ports play the roles of exporter and importer at the same time, the total ratio is around 80 percent if no repositioning is done.

Online linear programming (LP). In this method, the ECR problem is modeled as an integer linear programming (ILP) problem




Figure 4: Order proportion distribution between two ports in the original topology (a) and the new topology (b). Each column represents the proportions of orders from one port to orders. Note that two merged ports are the 9th and 10th columns.

with mathematical definitions. Here, we use the baseline method in [13], which applies the rolling horizon policy to do periodic planning. It solves the ILP problem for a long period and only adopts the plans of a short period, which prevents the deviation of planning from reality. The planning results are largely influenced by the unpredictable noise in the environment, which is the reason why it underperforms our method. Besides, the online computational cost of searching ILP solutions is much larger than our method.

DGN in A2C version. We chose one of the recent GNN based MARL frameworks called DGN as another baseline and adapt their key contributions into the multi-agent A2C framework, i.e., the network structure and temporal relation regularization. DGN directly uses the local reward to train, under the latent assumption that the local reward encourages cooperation, which is not always held in many real-world problems such as ECR. Therefore, we train it with artificially designed local reward, which is a combination of the global and the local reward. As you can see, we also outperform them in both normal and hard mode.

5.5 Ablation Study

The ablation studies are used to answer two questions: (1). Does the encoder-decoder attention in EncGAT help heterogeneous information aggregation and intricate interaction understanding? (2). Does PreLAC improve learning performance of cooperative policy?

Encoder-Decoder Attention. We compare our encoder-decoder attention with alternative aggregation functions, i.e. single decoder attention, GAT [30] and LSTM [7]. The single decoder attention directly puts the neighborhood sequence to the attention function in the decoder network introduced in Section 4.1. In the task of the normal mode, its performance is close to our method, but in the task of hard mode, the performance drops by 2%. The reason is that in the hard mode, the increased noise in the simulation data like the vessels' sailing time and the quantity of orders requires agents to cooperate against uncertainty, making the interactions more complicated. GAT is different from the single decoder attention in the computation of the attention coefficients. In the original implementation of GAT, the attention coefficients are calculated with a single-layer feedforward network. The result underperforms the dot-product based single decoder attention, which indicates that the way to calculate the attention coefficients may significantly influence the result. We also try the LSTM aggregator introduced in [7]. Although LSTM also has the ability to consider the inter-correlated influence from neighborhood sequence on an agent, the

Table 1: Fulfillment ratios (%) of different baselines and ablation experiments.

Section		Implementation Details	Simulation Mode	
			Normal	Hard
Baseline		No repositioning	81.95	80.34
		Online LP	94.00	87.99
		DGN in A2C	90.35 \pm 1.09	86.50 \pm 0.62
Ablation Study	PreLAC	Normal GC	85.44 \pm 0.10	80.77 \pm 0.62
		Separate actors normal GC	91.22 \pm 0.12	81.68 \pm 0.05
		Separate actors PreLAC	95.17 \pm 0.29	90.13 \pm 0.40
	EncGAT	Single decoder attention	98.32 \pm 0.02	90.44 \pm 0.54
		GAT	92.91 \pm 0.68	86.01 \pm 0.09
		LSTM	86.11 \pm 1.88	80.80 \pm 0.03
Our Method (EncGAT-PreLAC)			98.78 \pm 0.21	92.92 \pm 0.38

result is relatively poor. We think the reason is that the recurrent structure makes it difficult to learn with weak signals.


Pre-training with local actor-critic. To analyze how the pre-training algorithm improves learning performance, we first train the model directly with the global reward (Normal GC). It yields a poor performance especially in the hard mode, which is nearly the same as no repositioning at all. To further understand the result, we visualize the learned embedding of each port by mapping the output vector after EncGAT to the two-dimension space with principal component analysis (PCA). As is shown in Figure 5 where the data points of the same port are painted in the same color, we see that with the enhancement of pre-training, the data points regularly gather to several clusters compared to the mass without pre-training in the left picture. This indicates that without the local actor-critic task, the graph layers fail to learn distinguishable vertex embeddings, which will puzzle the succeeding actor and critic headers. An interesting experiment is also conducted to see the performance if we do not share the actor header between agents (separate actors normal GC), which leads to a little improvement. The reason is that the separate actor headers learn simple but different policies for agents regardless of the poor embedding input. Finally, we also try not to share actor headers in our framework (separate actors PreLAC). The result shows about a 2% drop in performance, which indicates that the parameter sharing, as a space to share experience between agents, helps train a better policy.

Inductiveness and generality. The inductiveness in the field of GNNs means the ability to generalize to unseen graph structures. Because the EncGAT itself is inductive and we share the actor headers among agents, the policy should be *inductive* too. Henceforth, we further conduct experiments to analyze the generality of the learned policy when the interaction graph changes, e.g., the changes of the topology and the trading distribution in the ECR problem. In particular, we first merge two ports into one port that takes the roles of the original two ports in their own route. Then, we randomly generate the trading distribution, which is different from the distribution generated from the WTO’s data. We plot the heat-map based on the trading quantity in the right of Figure 4 and compare it with the original one in the left of Figure 4. As the distribution completely differs from the original one, the role each port plays and the interactions between agents are different. The normal and hard modes are also defined in this setting. We test the performance of the model trained on the original setting on

Table 2: Fulfillment ratios (%) of the trained model and no repositioning baseline on two test topologies.

Mode	Trained model	No repositioning
Normal	78.69	67.77
Hard	76.01	65.10

this new topology *directly*. The result in Table 2 shows significant improvements in the fulfillment ratio compared to that with no repositioning.

**Figure 5: Visualization of vertex embeddings in global-critic (a) and our method (b). The color identifies different ports and each point represents a feature vector in a batch.**

6 CONCLUSION AND FUTURE WORK

We have proposed a GNN based MARL framework, EncGAT-PreLAC, to solve the emerging challenges in practical cooperative MAS. To extract cooperative information from the heterogeneous interaction graph, we leverage the encoder-decoder structure in graph attention to embed heterogeneous features as well as model intricate interactions. To alleviate the training difficulty of GNN under the weak global reward signal, we introduce the pre-training task empowered by the vertex-level local reward. Extensive experiments are conducted to solve a practical problem called ECR and the results show the advantages of our method. Nevertheless, some other challenges remain untouched. Firstly, the credit assignment problem becomes serious in larger-scaled cooperative tasks. It is interesting to combine techniques like difference rewards with our EncGAT for a better assignment. Secondly, the interaction information is sometimes unreliable because of either noise or deliberate interference like hacking in real-world problems. How to make cooperation robust to dirty information is practical and promising. We will conduct our work in these two aspects in the future.

REFERENCES

- [1] Akshat Agarwal, Sumit Kumar, and Katia Sycara. 2020. Learning transferable cooperative behavior in multi-agent team. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems*.
- [2] François-Xavier Devailly, Denis Larocque, and Laurent Charlin. 2020. IG-RL: Inductive Graph Reinforcement Learning for Massive-Scale Traffic Signal Control. arXiv:2003.05738 [cs.LG]
- [3] Rafael Epstein, Andres Neely, Andres Weintraub, Fernando Valenzuela, Sergio Hurtado, Guillermo Gonzalez, Alex Beiza, Mauricio Naveas, Florencio Infante, Fernando Alarcon, et al. 2012. A strategic empty container logistics optimization in a major shipping company. *Interfaces* 42, 1 (2012), 5–16.
- [4] Daniel J Fagnant and Kara M Kockelman. 2018. Dynamic ride-sharing and fleet sizing for a system of shared autonomous vehicles in Austin, Texas. *Transportation* 45, 1 (2018), 143–158.
- [5] Jakob Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. 2018. Counterfactual Multi-Agent Policy Gradients. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- [6] Supriyo Ghosh, Pradeep Varakantham, Yossiri Adulyasak, and Patrick Jaillet. 2017. Dynamic repositioning to reduce lost demand in bike sharing systems. *Journal of Artificial Intelligence Research* 58 (2017), 387–430.
- [7] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive Representation Learning on Large Graphs. In *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.). Curran Associates, Inc., 1024–1034.
- [8] Junling Hu, Michael P Wellman, et al. 1998. Multiagent reinforcement learning: theoretical framework and an algorithm.. In *Proceedings of the 25th International Conference on Machine Learning*, Vol. 98. Citeseer, 242–250.
- [9] Weihua Hu*, Bowen Liu*, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay Pande, and Jure Leskovec. 2020. Strategies for Pre-training Graph Neural Networks. In *Proceedings of the International Conference on Learning Representations*.
- [10] Arthur Jiang, Jia Zhang, Pingchao Yu, and etc. 2020. MARO: A Multi-Agent Resource Optimization Platform. <https://github.com/microsoft/maro>
- [11] Jiechuan Jiang, Chen Dun, and Zongqing Lu. 2020. Graph convolutional reinforcement learning for multi-agent cooperation. In *Proceedings of the International Conference on Learning Representations*.
- [12] Thomas N Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *Proceedings of the International Conference on Learning Representations*.
- [13] Xihan Li, Jia Zhang, Jiang Bian, Yunhai Tong, and Tie-Yan Liu. 2019. A cooperative multi-agent reinforcement learning framework for resource balancing in complex logistics network. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 980–988.
- [14] Yuan Li, Xiaodan Liang, Zhiting Hu, Yinbo Chen, and Eric P. Xing. 2019. Graph Transformer. <https://openreview.net/forum?id=HJei-2RcK7>
- [15] Yexin Li, Yu Zheng, and Qiang Yang. 2018. Dynamic bike reposition: A spatio-temporal reinforcement learning approach. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1724–1733.
- [16] Kaixiang Lin, Renyu Zhao, Zhe Xu, and Jiayu Zhou. 2018. Efficient large-scale fleet management via multi-agent deep reinforcement learning. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1774–1783.
- [17] Ning Liu, Zhe Li, Jielong Xu, Zhiyuan Xu, Sheng Lin, Qinru Qiu, Jian Tang, and Yanzhi Wang. 2017. A hierarchical framework of cloud resource allocation and power management using deep reinforcement learning. In *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 372–382.
- [18] Aleksandra Malysheva, Daniel Kudenko, and Aleksei Shpilman. 2019. MAGNet: Multi-agent Graph Network for Deep Multi-agent Reinforcement Learning. In *2019 XVI International Symposium "Problems of Redundancy in Information and Control Systems" (REDUNDANCY)*. IEEE, 171–176.
- [19] Kunal Menda, Yi-Chun Chen, Justin Grana, James W. Bono, Brendan D. Tracey, Mykel J. Kochenderfer, and David Wolpert. 2019. Deep Reinforcement Learning for Event-Driven Multi-Agent Decision Processes. *IEEE Transactions on Intelligent Transportation Systems* 20, 4 (Apr 2019), 1259–1268. <https://doi.org/10.1109/its.2018.2848264>
- [20] Nicolò Navarin, Dinh V. Tran, and Alessandro Sperduti. 2018. Pre-training Graph Neural Networks with Kernels. arXiv:1811.06930 [cs.LG]
- [21] Ling Pan, Qingpeng Cai, Zhixuan Fang, Pingzhong Tang, and Longbo Huang. 2019. A deep reinforcement learning framework for rebalancing dockless bike sharing systems. In *Proceedings of the AAAI conference on Artificial Intelligence*, Vol. 33. 1393–1400.
- [22] Liviu Panaft and Sean Luke. 2005. Cooperative Multi-Agent Learning: The State of the Art. *Autonomous Agents and Multi-Agent Systems* 11, 3 (2005), 387–434. <https://doi.org/10.1007/s10458-005-2631-2>
- [23] Tabish Rashid, Mikayel Samvelyan, Christian Schroeder, Gregory Farquhar, Jakob 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*. 4295–4304.
- [24] Dong-Ping Song and Jonathan Carter. 2009. Empty container repositioning in liner shipping. *Maritime Policy & Management* 36, 4 (2009), 291–307.
- [25] Dong-Ping Song and Jing-Xin Dong. 2012. Cargo routing and empty container repositioning in multiple shipping service routes. *Transportation Research Part B: Methodological* 46, 10 (2012), 1556 – 1575. <https://doi.org/10.1016/j.trb.2012.08.003>
- [26] Chuangchuan Sun, Macheng Shen, and Jonathan P How. 2020. Scaling Up Multiagent Reinforcement Learning for Robotic Systems: Learn an Adaptive Sparse Communication Graph. arXiv preprint arXiv:2003.01040 (2020).
- [27] Yizhou Sun and Jiawei Han. 2013. Mining Heterogeneous Information Networks: A Structural Analysis Approach. *SIGKDD Explor. Newsl.* 14, 2 (April 2013), 20–28. <https://doi.org/10.1145/2481244.2481248>
- [28] Peter Sunehag, Guy Lever, Audrunas Gruslys, Wojciech Marian Czarnecki, Vinicius Zambaldi, Max Jaderberg, Marc Lanctot, Nicolas Sonnerat, Joel Z Leibo, Karl Tuyls, et al. 2018. Value-decomposition networks for cooperative multi-agent learning based on team reward. In *Proceedings of the 17th international conference on autonomous agents and multiagent systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2085–2087.
- [29] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.). Curran Associates, Inc., 5998–6008.
- [30] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2018. Graph attention networks. In *Proceedings of the International Conference on Learning Representations*.
- [31] Xiao Wang, Houye Ji, Chuan Shi, Bai Wang, Yanfang Ye, Peng Cui, and Philip S Yu. 2019. Heterogeneous graph attention network. In *The World Wide Web Conference. 2022–2032*.
- [32] Hua Wei, Nan Xu, Huichu Zhang, Guanjie Zheng, Xinshi Zang, Chacha Chen, Weinan Zhang, Yanmin Zhu, Kai Xu, and Zhenhui Li. 2019. Colight: Learning network-level cooperation for traffic signal control. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 1913–1922.
- [33] Jun Wu, Xin Xu, Pengcheng Zhang, and Chunming Liu. 2011. A novel multi-agent reinforcement learning approach for job scheduling in grid computing. *Future Generation Computer Systems* 27, 5 (2011), 430–439.