

# Towards Generalizability of Multi-Agent Reinforcement Learning in Graphs with Recurrent Message Passing

Jannis Weil

Technical University of Darmstadt  
Darmstadt, Germany  
jannis.weil@tu-darmstadt.de

Osama Abboud

Huawei Technologies  
Munich, Germany  
osama.abboud@huawei.com

Zhenghua Bao

Technical University of Darmstadt  
Darmstadt, Germany  
zhenghua.bao@stud.tu-darmstadt.de

Tobias Meuser

Technical University of Darmstadt  
Darmstadt, Germany  
tobias.meuser@tu-darmstadt.de

## ABSTRACT

Graph-based environments pose unique challenges to multi-agent reinforcement learning. In decentralized approaches, agents operate within a given graph and make decisions based on partial or outdated observations. The size of the observed neighborhood limits the generalizability to different graphs and affects the reactivity of agents, the quality of the selected actions, and the communication overhead. This work focuses on generalizability and resolves the trade-off in observed neighborhood size with a continuous information flow in the whole graph. We propose a recurrent message-passing model that iterates with the environment’s steps and allows nodes to create a global representation of the graph by exchanging messages with their neighbors. Agents receive the resulting learned graph observations based on their location in the graph. Our approach can be used in a decentralized manner at runtime and in combination with a reinforcement learning algorithm of choice. We evaluate our method across 1000 diverse graphs in the context of routing in communication networks and find that it enables agents to generalize and adapt to changes in the graph.

## KEYWORDS

Multi-Agent Reinforcement Learning; Graph Neural Networks; Communication Networks

### ACM Reference Format:

Jannis Weil, Zhenghua Bao, Osama Abboud, and Tobias Meuser. 2024. Towards Generalizability of Multi-Agent Reinforcement Learning in Graphs with Recurrent Message Passing. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, Auckland, New Zealand, May 6 – 10, 2024, IFAAMAS, 9 pages.

## 1 INTRODUCTION

The capability of an adaptive system depends on the quality of its input. Ideally, it has access to the state and makes fully informed decisions at all times. Research in multi-agent reinforcement learning ranges from centralized to decentralized approaches [22]. Our focus

lies on decentralized systems in graph-based environments. While agents may leverage the complete state during training, they have limited access to local information during execution. In decentralized approaches, agents can directly react to local changes. However, having only local information may lead to suboptimal decisions. A common way to counteract this is to expand the observations of each agent by information from their direct neighborhood [30]. Including more nodes in the observed neighborhood improves decision making [5] but increases the communication overhead.

Recent works show that graph neural networks [34] and neural message passing [11] are well suited for applications in graph-based environments, especially because they can generalize to unseen graphs [33]. However, the approaches often assume a centralized view [1], explicit coordination across all agents [2], or the availability of labeled data in order to apply supervised learning [10].

We aim to resolve the trade-off in limited observed neighborhoods in graph-based environments and propose a recurrent approach where nodes exchange local information via message passing to improve their understanding of the global state. They refine their local states over the environment’s steps, allowing information to iteratively travel through the whole graph. Based on these node states, agents receive location-dependent graph observations.

Our approach provides a novel foundation for learning communication systems in multi-agent reinforcement learning and is jointly trained in an end-to-end fashion. Our contributions are:

- We propose to *decouple* learning graph-specific representations and control by separating node and agent observations.
- To the best of our knowledge, we are the first to address the problem of limited observed neighborhoods in graph-based environments with recurrent graph neural networks.
- We show that the learned graph observations enable generalization over 1000 diverse graphs in a routing environment, achieving similar throughput as agents that specialize on single graphs when combined with action masking.
- We show that our approach enables agents to adapt to a change in the graph on the fly without retraining.

Our code is available at <https://github.com/jw3il/graph-marl>. The remainder of this paper is structured as follows. We begin with the problem statement in Sec. 2 and then introduce our approach in Sec. 3. We describe our evaluation setup in Sec. 4, the results are presented and discussed in Sec. 5. The following Sec. 6 provides an overview of related work and Sec. 7 concludes the paper.



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

*Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems ([www.ifaamas.org](http://www.ifaamas.org)).

## 2 REINFORCEMENT LEARNING IN GRAPHS

Reinforcement learning (RL) in graph-based environments has gained much popularity in many application domains [28], including communication networks [22]. We consider multi-agent environments that build upon a graph  $G \doteq (V, E) \in \mathcal{G}$ , representing a communication network of nodes  $V$  connected via undirected edges  $E \subseteq V \times V$ . We assume that the agents  $I$  are part of a partially observable stochastic game [16] that requires them to consider the graph, i.e. state  $s_t \in S$  at step  $t$  contains  $G$  and may augment it with state information that characterizes the network. Examples include edge delays and compute resources of nodes. Each agent  $i \in I$  receives partial observations  $o_t^i \sim O^i(\cdot | s_t)$  and selects an action  $a_t^i \in A$  using its policy  $a_t^i \sim \pi^i(\cdot | o_t^i)$ . The agent’s goal is to maximize its expected discounted return  $\mathbb{E}[\sum_{t'=t}^T \gamma^{t'-t} R_{t'}^i]$  with discount factor  $\gamma \in [0, 1]$  and time horizon  $T$ , where the individual rewards  $R_t^i = R^i(s_t, a_t)$  are based on the joint action of all agents  $a_t = (a_t^i)_{i \in I}$ .

How agents observe the graph and the network’s state is usually not discussed in depth by related works. A centralized view allows for the best decision making but does not scale to bigger graphs. A decentralized view trades off reactivity and amount of available information with the size of the observed neighborhood. Authors usually decide for one of these views and their approaches are therefore, by design, limited to certain problems or graph structures.

Consider a decentralized approach where an agent is located on a node and observes its  $n$ -hop neighborhood. Depending on the environment, the assumption about the observable neighborhood is critical with respect to generalizability. Let’s imagine the agent is supposed to find the shortest path to some destination node in the graph that is  $n + 1$  hops away. How can it identify the shortest path to a destination node that’s not included in its observation? There are two straight-forward solutions:

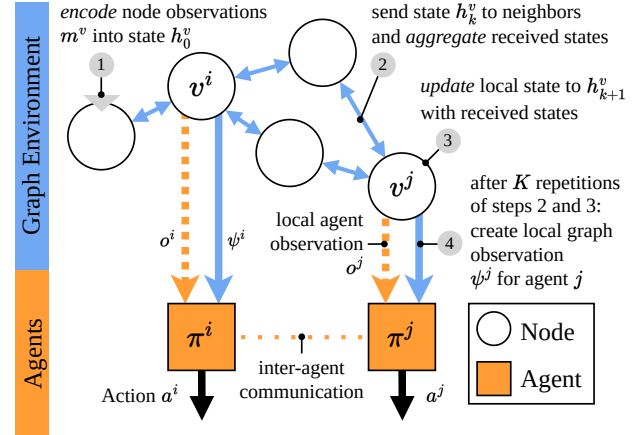
- (1) Specialization on a single graph. Agents explore the whole graph during training. If the graph is fixed and nodes are uniquely identifiable based on the agent’s observation, agents can specialize and find the optimal path to any node.
- (2) Expansion of the observation space to include the missing information, e.g. to  $(n + 1)$ -hop neighborhoods.

With specialization, the learned solution will not generalize to other graphs. If the target graph does not change, this would be a sufficient but highly inflexible solution. As real networks have diverse underlying graphs and are usually dynamic, many researchers consider online training on the target graph. However, as reinforcement learning requires agents to explore and make suboptimal decisions, online training from scratch might be unacceptable in practice. In contrast, expanding the observation space allows agents to generalize. However, the observation range to consider greatly depends on the concrete problem. For example, for routing, global knowledge is necessary if any node could be the destination. But then the approach is not decentralized anymore and will not scale.

Our idea is to address this issue by expanding the observation space with learned graph observations that leverage recurrent message passing. Agents are still reactive and don’t have to gather information about the whole graph before making a decision. While initial decisions may be suboptimal, the quality of the learned graph observations should increase over time. Ideally, they converge to a global view and allow agents to make optimal decisions.

## 3 LEARNED GRAPH OBSERVATIONS

We consider environments that are based on a graph and propose to *decouple* both components, i.e. agents do not have to keep track of the whole graph state and can build upon a lower-level mechanism that aggregates graph information. The information flow is illustrated in Fig. 1 and will be explained in the following subsections.



**Figure 1: Our graph observation mechanism iteratively distributes node states via message passing. Agents in the graph receive local graph observations for decision making.**

### 3.1 Recurrent Message Passing

The core idea of our graph observations is leveraging the message-passing framework of graph neural networks [13] to distribute local information in the network. Related approaches are usually used in a centralized manner with a global view [3]. However, recurrent aggregation functions [23, 37] spread multiple message passing iterations over time and enable decentralization [10]. We introduce a second recurrent loop back to the input of the graph neural network and label this approach *recurrent message passing*.

We assume that each node  $v \in V$  receives local node observations  $m^v \sim M^v(\cdot | s)$  based on an unknown system state  $s \in S$ . Instead of directly using this as an input of a graph neural network, each node  $v \in V$  embeds its local observation into its current node state  $h^v$  with an arbitrary differentiable function *encode*:

$$h_0^v \doteq \text{encode}(h^v, m^v). \quad (1)$$

The node state  $h^v$  is initialized to zero in the first step and will serve as a recurrent loop between subsequent environment steps. This allows nodes to consider previously aggregated information when processing observations, similar to auto-regressive graph models [31] that use predictions of previous steps as their input.

At iteration  $k \geq 0$ , each node sends their state  $h_k^v$  to all direct neighbors  $w \in N(v) \doteq \{w | (v, w) \in E\}$ . Each neighbor generates a new state by first aggregating incoming node states and then updating its state using arbitrary differentiable functions *aggregate*<sup>k</sup> and *update*<sup>k</sup>. One message passing iteration is defined as:

$$M_k^v \doteq \text{aggregate}^k((h_k^w)_{w \in N(v)}) \quad (2)$$

$$h_{k+1}^v \doteq \text{update}^k(h_k^v, M_k^v). \quad (3)$$

**Algorithm 1** Distributed Node State Update

**Input:** Node  $v$  with direct neighbors  $N(v)$ , state  $s$ , previous node state  $h^v$ , node observation  $m^v$

**Output:** Updated node state  $h_K^v$  and intermediate states  $H^v$

- 1:  $h_0^v \leftarrow \text{encode}(h^v, m^v)$  ▷ Encode node observation
- 2: **for**  $k \leftarrow 0$  to  $K - 1$  **do** ▷ Update with message passing
- 3:   Send  $h_k^v$  to all neighbors  $w \in N(v)$
- 4:   Receive  $h_k^w$  from all neighbors  $w \in N(v)$
- 5:    $M_k^v \leftarrow \text{aggregate}^k((h_k^w)_{w \in N(v)})$
- 6:    $h_{k+1}^v \leftarrow \text{update}^k(h_k^v, M_k^v)$
- 7:  $H^v \leftarrow (h_k^v)_k \parallel (h_k^w)_{w \in N(v), k}$  ▷ Get all intermediate states
- 8: **return**  $h_K^v, H^v$

Alg. 1 shows pseudo code for recurrent message passing that is executed by all nodes  $v \in V$ . There are  $K \in \mathbb{N}$  iterations between steps in the environment, i.e. equations (2) and (3) are repeatedly applied until  $k + 1 = K$ . The final aggregate  $h_K^v$  and all intermediate node states  $H^v$  received and calculated by  $v$  can then be used by agents, as we will detail later. In the next environment step, we set  $h^v = h_K^v$  and repeat the algorithm. The number of iterations can be adjusted to fit the requirements of the learning task. An increased number of iterations per step causes information to traverse the network faster but also increases the communication overhead.

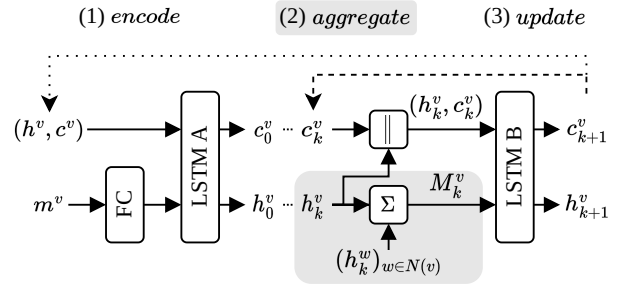
This node state update is performed at each step in the environment. We assume that each agent  $i \in I$  is assigned to exactly one node  $v^i \in V$  at each step. In addition to its observation in the environment, agent  $i \in I$  then receives a local graph observation  $\psi^i$  of the node it is assigned to based on the information  $H^{v^i}$  this node received in this step via a differentiable readout function  $\Psi$ :

$$\psi^i \doteq \Psi(H^{v^i}). \quad (4)$$

In the simplest case,  $\Psi$  could be the identity function of the latest node state  $\Psi(H^{v^i}) \doteq h_K^{v^i}$  in each iteration, i.e. agents receive the current state of the node to which they are assigned. Access to intermediate node states is necessary to allow for skip connections and aggregation mechanisms like jumping knowledge networks [42].

### 3.2 Model Architecture

Based on the design from the previous section, we propose a simple recurrent message passing architecture. The *encode* function is represented by a fully connected network to embed the node observation and an LSTM [17] to update the previous node state with the new embedding. The *aggregate* function is the sum of all neighbors' hidden states, but could be any graph convolution from related work. Finally, *update* is modeled by another LSTM. We share parameters for all iterations, i.e.  $\forall k. \text{update}^k = \text{update}$ . We provide an overview of the architecture with Fig. 2. An LSTM cell takes an input tensor and a pair of hidden and cell state tensors  $(h, c)$  and yields new hidden and cell states. In our architecture, the hidden state is exchanged with neighbor nodes during aggregation, the cell state remains local to the node. The inner loop from update to aggregate depicts iterations within a step, the outer loop represents the forwarding of states between environment steps. The states



**Figure 2: Our recurrent message passing model leverages LSTM cells to encode the node observation and update the node state. The hidden states of neighbor nodes are aggregated via summation, cell states remain local to each node.**

of the LSTM modules are not separated. Instead, a single pair of hidden and cell states is passed on between the modules.

### 3.3 Integration in Deep RL

Graph observations are compatible with all deep reinforcement learning approaches that allow for backpropagation through the policy's input, i.e. the observation space. To the best of our knowledge, most algorithms do not have any limitations in that regard, as the policy is usually based on a differentiable function.

In order to integrate our method with deep RL algorithms, the node state update has to be performed at each environment step during inference and training, resulting in an expanded observation space for the agents. The integration at inference time can be done with a simple environment wrapper. Depending on the considered algorithm, the integration into the training loop will require additional effort. For example, algorithms based on Q-learning bootstrap the value target using the observations of future states. In order to compute the corresponding graph observations, we therefore have to compute or sample node states for these future steps. Additionally, as node states are updated over multiple environment steps, we have to unroll the node state update over a sequence of steps and apply backpropagation through time in order to learn stable update functions. This is already included in algorithms that consider stateful agents with recurrent models [20], but will require adjustments for other reinforcement learning algorithms.

### 3.4 Exemplary Integration in Deep Q-Learning

In this section, we exemplarily describe how to integrate our method into independent DQN [27] with parameter sharing across agents. Our approach is summarized in Alg. 2, noteworthy changes to the original algorithm are highlighted in light gray. The main difference lies in the introduction of node states  $h_t$  that are updated in parallel to the environment steps based on the node observations  $m_t$  and the previous node states. For notational simplicity, we denote recurrent message passing combined with the readout function  $\Psi$  as a differentiable function  $U(h_t, m_t, s_t; \theta_U)$  parameterized by  $\theta_U$  (see line 9). It returns the next node state of all nodes  $h_{t+1} \doteq (h_{t+1}^v)_{v \in V}$  and the graph observations of all agents  $\psi_t \doteq (\psi_t^i)_{i \in I}$  based on the node states of all nodes  $h_t$ , all node observations  $m_t$  and the

state  $s_t$ . The only information required from  $s_t$  are the graph and the mapping of agents to nodes. The graph observation  $\psi_t^i$  of agent  $i$  depends on its position in the graph and is concatenated with the agent’s observation  $o_t^i$  received from the environment (see line 12).  $\hat{Q}$  and  $\hat{U}$  denote that the respective gradient calculation is disabled.

During training, we sample a sequence of transitions from the replay memory and perform backpropagation through time analogous to the stored state method from recurrent experience replay [20]. The initial node state  $h'_{j_0}$  is loaded from the replay memory and subsequent node states in the sequence are recomputed (see lines 18 and 19). The Q-learning target requires graph observations for the next step, which we compute temporarily. We then aggregate the squared error over all steps in the sequence (see line 22) and perform gradient descent with respect to the agent’s parameters  $\theta_Q$  and the parameters of the message passing module  $\theta_U$ .

---

**Algorithm 2** Independent DQN with Learned Graph Observations
 

---

```

1: Initialize replay memory  $D$ 
2: Initialize action-value function  $Q$  with weights  $\theta_Q$ 
3: Initialize target weights  $\hat{\theta}_Q$ 
4: Initialize node state update function  $U$  with weights  $\theta_U$ 
5: for episode  $\leftarrow 0 \dots \mathbf{do}$ 
6:    $h_0 \leftarrow 0$  ▷ Initialize node states
7:   Obtain  $s_0, o_0,$  and  $m_0$  by resetting the environment
8:   for  $t \leftarrow 0$  to  $T - 1$  do
9:      $h_{t+1}, \psi_t \leftarrow \hat{U}(h_t, m_t, s_t; \theta_U)$  ▷ Node state update
10:    for  $i \in I$  do
11:      Select random action  $a_t^i$  with probability  $\epsilon$ 
12:      otherwise select action
13:       $a_t^i = \operatorname{argmax}_a \hat{Q}(o_t^i \parallel \psi_t^i, a; \theta_Q)$ 
14:    Perform environment step with actions  $a_t$  and get
15:    reward  $r_t$ , state  $s_{t+1}$ , obs  $o_{t+1}$ , node obs  $m_{t+1}$ 
16:    Store  $(h_t, h_{t+1}, m_t, m_{t+1}, s_t, s_{t+1}, o_t, a_t, r_t, o_{t+1})$  in  $D$ 
17:    Initialize loss  $L \leftarrow 0$ 
18:    for batch sequence indices in  $D$   $j \leftarrow j_0$  to  $j_0 + (J - 1)$  do
19:      if  $j = j_0$  then
20:         $h'_j \leftarrow h_j$  ▷ Load node state from replay memory
21:         $h'_{j+1}, \psi'_j \leftarrow U(h'_j, m_j, s_j; \theta_U)$  ▷ Train node state update
22:         $h''_{j+2}, \psi''_{j+1} \leftarrow \hat{U}(h'_{j+1}, m_{j+1}, s_{j+1}; \theta_U)$  ▷ Target input
23:         $y_j \leftarrow r_j + \mathbf{Z}_{j\gamma} \max_a \hat{Q}(o_{j+1} \parallel \psi''_{j+1}, a; \hat{\theta}_Q)$ 
24:        with  $\mathbf{Z}_j^i = \begin{cases} 0 & \text{if agent } i \text{ is done at step } j + 1 \\ 1 & \text{otherwise} \end{cases}$ 
25:         $L \leftarrow L + (y_j - Q(o_j \parallel \psi'_j, a_j; \theta_Q))^2$ 
26:      Perform gradient descent on  $L$  with respect
27:      to parameters  $\theta_Q$  and  $\theta_U$ 
28:      Update target weights  $\hat{\theta}_Q$ 

```

---

## 4 EXPERIMENT SETUP

We evaluate our approach in diverse graphs based on a routing environment. The following sections describe considered models, algorithms and graphs with greater detail. Then we briefly describe the routing environment and a simplified supervised learning task.

### 4.1 Models and Training Algorithms

Our design consist of two parts, a model that generates graph observations and a reinforcement learning agent.

*Graph Observations.* The core of the graph observations is the message passing framework, as described in Sec. 3.1. Any graph neural network can be used to generate such graph observations. We use our proposed architecture from Sec. 3.2 and consider three baseline graph neural network architectures from related work with implementations by PyTorch Geometric [9] and PyTorch Geometric Temporal [32]. Two architectures are feed-forward graph neural networks without recurrency. GraphSAGE [14] is a GNN with multiple graph convolutional layers that use individual parameters. A-DGN [12] aims to improve learning long-range dependencies with an added diffusion term and performs multiple iterations with the same parameters. As a recurrent baseline, GCRN-LSTM [37] combines an LSTM with Chebyshev spectral graph convolutions [7]. While our architecture uses a single sum to aggregate hidden states, GCRN-LSTM utilizes 8 Chebyshev convolutions to aggregate intermediate computations of an LSTM cell.

We define the readout function  $\Psi$  of aggregated node state update information  $H^v$  (see Sec. 3.1) to graph observations  $\psi^i$  as a concatenation of the current node state  $h_K^v$  and the last node states  $h_{K-1}^w$  that this node received from its neighbors. This serves as a skip connection over the last iteration. Note that no additional message exchange is necessary for this skip connection. We apply the same readout function to all graph observation methods.

*Agents.* We consider independent DQN [27], recurrent DQN (DQNR) [20], CommNet [38] and DGN<sup>1</sup> [19]. We build upon the implementation of DGN<sup>2</sup> and reimplement the remaining approaches. All variants share the same training setup but differ in the agent’s architecture. DQN is a feed-forward network with fully-connected layers that is trained with a Q-learning loss. DQNR adds an LSTM [17] layer and is trained on sequences. Both approaches do not feature any information exchange between agents, their policies are completely separated during execution. CommNet extends DQNR with two communication rounds where agents exchange their hidden states before selecting an action. DGN extends DQN with two communication rounds using self-attention [39] and adds a regularization term to the loss. Within one communication round of CommNet and DGN, agents communicate with other agents that reside on the same node or on a node in their direct neighborhood.

### 4.2 Graph Generation and Overview

We extend the graph generation used in the routing environment from Jiang et al. [19]. It places  $L$  nodes randomly on a 2D plane and then connects close nodes with edges until all nodes reach degree  $D$ . Having a fixed node degree is not a mandatory constraint for our approach, but results in a discrete action space of fixed size that simplifies reinforcement learning. Technically, this can be extended to graphs with nodes of variable degrees, e.g. via action masking [36]. The delay of an edge in steps is determined by a linear function of the distance between the connected nodes, rounded to the next integer. Disconnected graphs are filtered out.

<sup>1</sup>Not to be confused with the graph neural network A-DGN.

<sup>2</sup><https://github.com/PKU-RL/DGN/>, including the PyTorch version.

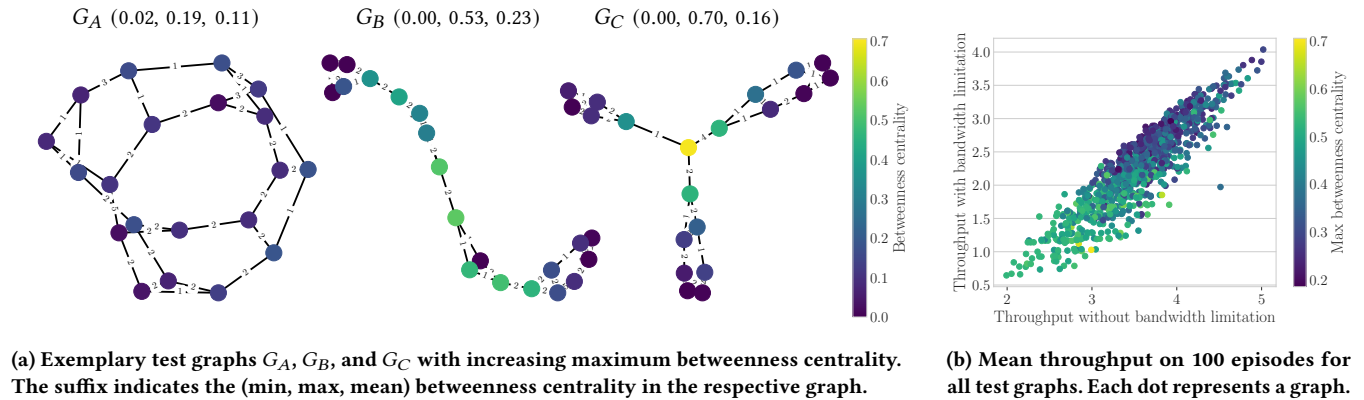


Figure 3: Overview of the considered graphs with (a) three exemplary graphs from the test set and (b) the mean throughput of shortest paths routing with and without bandwidth limitation in all 1000 test graphs.

We generate 1000 distinct graphs for testing with  $L \doteq 20$  and  $D \doteq 3$ . The mean diameter is  $7.21 \pm 1.42$  hops and  $12.84 \pm 2.72$  steps. The mean all-pairs shortest paths (APSP) lengths are  $3.26 \pm 1.92$  hops and  $5.7 \pm 3.6$  steps. The maximum APSP lengths equal the max diameters of 12 hops and 23 steps. The betweenness centrality in  $[0, 1]$  of a node reflects the proportion of shortest paths between any two different nodes in the graph that contain this node. In routing, high values indicate potential bottlenecks in the graph. We show exemplary graphs with increasing maximum betweenness centrality in Fig. 3a. The nodes are repositioned to provide a better overview. Graph  $G_A$  with a low maximum betweenness centrality is well balanced. Graph  $G_B$  has a high mean betweenness centrality due to its line-like structure. Graph  $G_C$  has a high maximum betweenness centrality because of the bottleneck node in the center.

Apart from the node connectivity and potential bottlenecks, the diameter of the graphs and the distribution of the shortest paths are expected to influence our approach. In the graph neural network architecture considered in this paper, messages only traverse the graph through its edges. The number of iterations for information from node  $v_1$  to be forwarded to node  $v_2$  equals the length of the shortest path between these nodes. The minimum number of iterations required to collect information from all nodes is therefore the diameter of the graph. While the maximum diameter is 12 hops, we found that over 99% of the shortest paths in all test graphs have at most 8 hops. Further details are provided in the appendix [41].

### 4.3 Routing Environment

We extend the routing environment from Jiang et al. [19] and fix a bug that caused packets to skip edges. At all times, there are  $N$  packets of random sizes in  $[0, 1)$  that have to be routed from random source to destination nodes in a given graph. We focus on generalizability across graphs and use  $N \doteq 20$  packets in our experiments. Each packet is an agent that receives a reward of 10 when it reaches its destination. On a node, agents select one of  $1 + D$  discrete actions that correspond to waiting and choosing an outgoing edge. Each edge has a transmission delay given in steps. We consider two environment modes. The first mode has no restrictions and packets are always transmitted via their selected edges. In the second mode, we take packet sizes and limited edge capacities into account. A packet

of size  $g$  is transmitted via a selected edge if the cumulative size of all packets that are currently transmitted via this edge is smaller than  $1 - g$ . The packet then traverses the edge according to the number of steps in its transmission delay. Otherwise, the packet is forced to stay at its current position and receives a penalty of  $-0.2$ . An agent’s local observations include its current position, its destination and packet size. For each outgoing edge of their current node, it observes the delay, the cumulative size of packets on that edge and the respective neighbor’s node id. A node observes its own id, the number and size of packets that reside on the node, and local information about outgoing edges. All node ids are given as one-hot encodings.

The *throughput* refers to the number of packets per step that arrive at their destination. *Delay* describes the length of their episodes. Note that the delay should never be considered on its own. For example, agents that only route to destinations in their 1-hop neighborhood would achieve low delays but also a low throughput.

As a baseline, we consider heuristic agents with a global view that always choose the shortest paths with respect to the edge delays. Fig. 3b shows the throughput with and without bandwidth limitations when using this heuristic for 100 episodes in all test graphs. Each dot is colored according to the maximum betweenness centrality of the respective graph. We can see that bandwidth limitations cause a significant drop in throughput and that graphs with high maximum betweenness tend to result in lower throughput.

### 4.4 Shortest Paths Regression Task

The routing environment requires agents to learn paths from source to destination nodes. To quickly evaluate the efficacy of different graph neural network architectures, we design a multi-target regression problem as a simplification of the routing environment. We expect that the performance of different architectures in this task will indicate their suitability for the routing environment. The training dataset contains node observations for 100 000 graphs generated by resetting the routing environment. We exclude 1 000 of these graphs for validation. The targets for each node are the shortest path lengths to all other nodes. For the test dataset, we use the node observations and targets of the 1 000 graphs from Sec. 4.2. The loss is the mean squared error between the predicted and real distances for each source and destination node.

## 5 RESULTS

We first present independent results for our two core components, the graph neural network architectures (see Sec. 5.1) and agents trained in the routing environment with single graphs (see Sec. 5.2). Section 5.3 combines both components and provides the results for generalized routing, followed by a discussion in Sec. 5.4.

We use the AdamW optimizer [25] for all experiments. Details regarding the hyperparameters are provided in the appendix [41].

### 5.1 Shortest Paths Regression

We first evaluate the considered graph neural network architectures in the shortest paths regression task (see Sec. 4.4). We train each architecture with three seeds for 50 000 iterations of batch size 32.


Based on the observations regarding the APSP distribution from Sec. 4.2,  $K = 8$  message passing iterations allow to pass information between over 99% of all node pairs in the test graphs. Therefore, we hypothesize that  $K = 8$  should lead to good performance on the test graphs for the non-recurrent models. The results for different message passing iterations  $K$  and unroll depths  $J$  are shown in Tab. 1. In GraphSAGE,  $K$  refers to the number of graph convolutional layers. For GCRN-LSTM, we set the filter size of the Chebyshev convolutions to  $K + 1$ . Both result in  $K$  message passing iterations. All approaches learn to approximate the shortest path lengths and achieve a mean squared error (MSE) of around or below 1 for at least one configuration. In the case of GraphSAGE, increasing the number of layers to 16 leads to unstable training and a high test loss in this task. As the non-recurrent architectures are stateless, they yield the same results at each forward step  $t$ .

For the recurrent approaches, we want to use a low number of iterations per step (i.e.  $K = 1$ ) to reduce the communication overhead. We evaluate them with different unroll depths  $J$  and higher values of  $K$  for comparison. As expected, the recurrent approaches perform poorly in the first forward step with  $t = 1$ . They refine their hidden states in subsequent steps and approximately reach their minimum losses at the unroll depth  $J$  that was used during training. Afterwards, we can see that their losses increase again. Increasing the unroll depth improves long-term stability, but leads to increased training time. A higher number of iterations  $K$  predominantly leads to improved predictions, at the cost of increasing the communication overhead per step.

For our experiments in combination with reinforcement learning, we select  $K = 8$  for the non-recurrent models and  $K = 1$ ,  $J = 8$  for the recurrent models. This is a compromise between performance, stability and communication overhead. Fig. 4 shows the validation loss of the selected approaches during training. The recurrent architectures converge faster to a low loss value than the non-recurrent ones. This could possibly be caused by better gradients, as we compute separate losses for each message passing iteration when unrolling the network. Both recurrent approaches achieve similar validation losses, although our architecture is simpler and exchanges less information during the forward steps. The high standard deviation of GCRN-LSTM in the beginning is caused by one of the three runs, where the validation loss does not decrease initially. In the reinforcement learning setting, we expect recurrent experience replay with stored states to further improve the long-term stability of the recurrent approaches.

**Table 1: Results for the shortest paths regression problem.  $K$  denotes the number of message passing iterations and  $J$  the unroll depth for recurrent approaches. Shown is the MSE on all test topologies after  $t$  forward steps with  $K$  iterations. All results are averaged over 3 seeds.**

Architecture	$K$	$J$	MSE at Forward Step $t$					
			1	2	4	8	16	32
GraphSAGE	8	-	1.16 (all seeds: 1.14, 1.22, 1.13)					
	16	-	3.57 (all seeds: 4.18, 3.46, 3.06)					
A-DGN	8	-	1.50 (all seeds: 1.49, 1.56, 1.46)					
	16	-	1.18 (all seeds: 1.16, 1.20, 1.18)					
GCRN-LSTM	1	8	4.98	2.98	1.12	<b>0.60</b>	1.61	4.27
	1	16	5.03	3.09	1.28	0.60	<b>0.53</b>	1.08
	2	8	3.18	1.22	0.49	<b>0.40</b>	0.51	1.08
	4	8	1.56	0.69	0.45	<b>0.43</b>	0.52	1.00
Ours	1	8	4.98	2.91	0.94	<b>0.43</b>	0.99	3.75
	1	16	5.02	2.98	1.02	0.39	<b>0.37</b>	0.46
	2	8	3.02	1.07	0.39	<b>0.35</b>	0.57	1.78
	4	8	1.26	0.48	<b>0.34</b>	<b>0.34</b>	0.42	0.81



**Figure 4: Validation loss of the selected GNN architectures in the shortest paths problem during training. The shaded area shows the standard deviation over 3 models.**

### 5.2 Routing in Single Graphs

Before we evaluate our method on multiple graphs, we train agents without graph observations in the routing environment using single graphs. The agents are trained for 250 000 total steps with 24 000 iterations of batch size 32. Episodes are truncated after 300 steps. In Tab. 2, we show the results for the outlier graph from Fig. 3b at around (4.5, 2.0). The top half shows the results without bandwidth limitations, the bottom half with bandwidth limitations. Without limitations, all methods learn the optimal shortest paths. This is not surprising, as the graph is static and agents can locate themselves in the graph using the node id they receive in their observation. The agents specialize on the graph. With bandwidth limitations, the throughput drastically decreases and the delay increases. We can see that in this graph, all learning approaches are able to outperform the shortest path solution in terms of mean reward and throughput. Except for DQN, the delay of arrived packets is slightly higher. Surprisingly, the effect of communication is very small. Using the same training setup for all agent architectures, we cannot reproduce the results from Jiang et al. [19] in this particular graph and find that the performance of DQN is very close to DGN.

**Table 2: Results for routing in a selected graph, averaged over 1000 episodes and 3 models. The learning approaches outperform the shortest paths heuristic.**

Mode	Agent	Metrics		
		Reward	Delay	Throughput
no limitation	ShortestPath	2.26 ± 0.0	4.39 ± 0.0	4.52 ± 0.0
	DQN	2.26 ± 0.0	4.39 ± 0.0	4.52 ± 0.0
	DQNR	2.26 ± 0.0	4.39 ± 0.0	4.52 ± 0.0
	CommNet	2.26 ± 0.0	4.39 ± 0.0	4.52 ± 0.0
	DGN	2.26 ± 0.0	4.39 ± 0.0	4.52 ± 0.0
bandwidth limitation	ShortestPath	0.88 ± 0.0	7.06 ± 0.01	1.98 ± 0.0
	DQN	1.05 ± 0.00	<b>6.81 ± 0.08</b>	2.15 ± 0.00
	DQNR	1.09 ± 0.00	7.20 ± 0.06	2.22 ± 0.01
	CommNet	1.11 ± 0.00	7.22 ± 0.03	<b>2.26 ± 0.01</b>
	DGN	1.10 ± 0.00	7.23 ± 0.13	2.25 ± 0.01

**Table 3: Throughput for routing agents individually trained on the graphs from Fig. 3a with varying betweenness centrality. Shown are the results for the bandwidth limitation mode.**

Agent	Throughput		
	Graph $G_A$	Graph $G_B$	Graph $G_C$
ShortestPath	3.20 ± 0.00	<b>1.53 ± 0.00</b>	<b>1.02 ± 0.00</b>
DQN	3.28 ± 0.01	1.43 ± 0.02	0.98 ± 0.01
DQNR	3.28 ± 0.00	1.47 ± 0.01	0.99 ± 0.01
CommNet	<b>3.31 ± 0.00</b>	<b>1.53 ± 0.00</b>	1.01 ± 0.01
DGN	3.29 ± 0.00	1.43 ± 0.05	1.00 ± 0.00

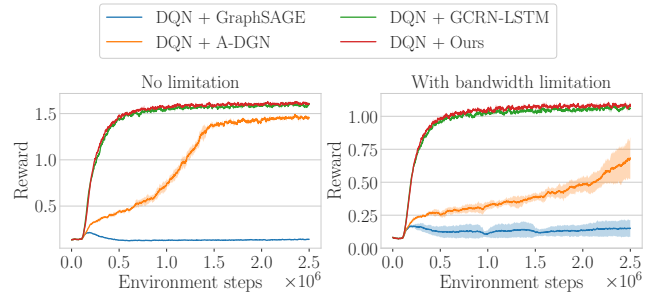
While we did not train agents for all 1000 test graphs, we have made similar observations for the other graphs we investigated. Tab. 3 shows the throughput in the three graphs from Fig. 3a, averaged over 1000 episodes and 3 models. We omit the results without limitations, as all approaches match shortest paths with a mean throughput of around 4.05 in graph  $G_A$ , 2.38 in graph  $G_B$  and 2.98 in graph  $G_C$ . For the limited bandwidth mode, we again find no notable difference between DQN and DGN in these graphs and notice that the learning approaches outperform shortest paths only for graph  $G_A$ . The throughput achieved by the learning approaches is approximately on par with shortest paths for graph  $G_B$  and  $G_C$ .

### 5.3 Generalized Routing

This section presents the results for our learned graph observations. We expect them to enable agents to generalize over different graphs. As all agent architectures achieve similar results for single graphs, we select DQN as the underlying agent architecture due to its simplicity. Instead of only receiving observations from the environment, agents now also receive graph observations from nodes they are located at. We train graph observations and agents end-to-end with reinforcement learning on randomly generated graphs for 2.5 million total steps, 240 000 iterations and batch size 32. Episodes are truncated after 50 steps to increase the number of generated graphs. The resulting models are evaluated on our 1000 test graphs and 300 episode steps for comparability with Sec. 5.2.

**Table 4: Results for routing in 1000 test graphs for 300 steps using DQN with graph observations provided by the listed methods. An asterisk (\*) indicates action masking at test time.**

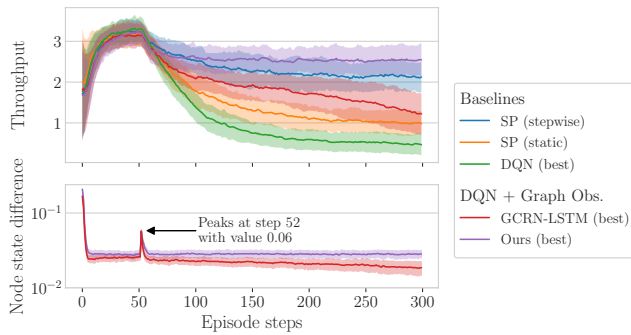
Mode	Method	Metrics		
		Reward	Delay	Throughput
no limitation	ShortestPath	<b>1.77 ± 0.00</b>	5.59 ± 0.00	<b>3.54 ± 0.00</b>
	GraphSAGE	0.02 ± 0.00	<b>4.49 ± 0.35</b>	0.04 ± 0.00
	A-DGN	1.15 ± 0.02	5.72 ± 0.02	2.29 ± 0.04
	GCRN-LSTM	1.55 ± 0.01	5.60 ± 0.01	3.09 ± 0.01
	Ours	1.56 ± 0.03	5.57 ± 0.02	3.12 ± 0.07
	Ours*	1.74 ± 0.00	5.65 ± 0.00	3.49 ± 0.00
bandwidth limitation	ShortestPath	1.05 ± 0.00	7.93 ± 0.00	2.26 ± 0.00
	GraphSAGE	0.02 ± 0.02	14.58 ± 6.85	0.08 ± 0.05
	A-DGN	0.30 ± 0.14	10.85 ± 1.34	0.67 ± 0.28
	GCRN-LSTM	0.99 ± 0.01	8.37 ± 0.02	2.03 ± 0.01
	Ours	1.02 ± 0.01	8.26 ± 0.02	2.10 ± 0.02
	Ours*	<b>1.10 ± 0.00</b>	<b>7.59 ± 0.01</b>	<b>2.38 ± 0.01</b>



**Figure 5: Reward of DQN with graph observations during training without (left) and with (right) bandwidth limitations. The shaded area shows the standard deviation over 3 models.**

Tab. 4 shows the average results over 3 seeds. The reward during training is shown in Fig. 5. The high positive reward and throughput of GCRN-LSTM and our proposed architecture show that graph observations indeed enable agents to generalize over different graphs. Our method achieves comparable results while having a lower communication overhead. However, we find that the results are worse than for agents that specialize on single graphs (see Sec. 5.2). For the non-recurrent approaches, A-DGN learns graph observations but converges much slower than the recurrent approaches. GraphSAGE fails to learn, even in experiments with batch size 256 and jumping knowledge networks [42] that are not shown here. Considering the results of Sec. 5.1, it is unclear why the non-recurrent approaches perform poorly in this setting. We hypothesize that the targets provided by backpropagation through time facilitate learning, but further experiments in different environments would be required to verify this. All methods have a lower throughput than the shortest paths heuristic without additional modifications.

Upon closer inspection of the behavior of a model trained with our architecture, we notice that around 12% of the 1000 test episodes contain packets that never arrive at their destination within 300 steps. This is caused by routing loops, a common issue that can be addressed with post-processing of the learned policy [18]. When



**Figure 6: Throughput over time and node state differences of selected models in graph  $G_A$  from Fig. 3a averaged over 100 episodes. The delay of a single edge is increased from 2 to 10 at step 50. The shaded area shows the standard deviation.**

repeating the experiment with different seeds, we observe routing loops in different graphs. We investigate an action masking mechanism that stores the path of a packet and masks actions that lead to already visited nodes. If there are no legal actions, the packet is dropped and a new packet spawns at a random location.

Action masking results in throughput improvements that match our expectations from the fixed topology setting, as shown in Tab. 4. However, it introduces 0.01 and 0.1 dropped packets per step for routing without and with bandwidth limitations, respectively.

#### 5.4 Adaptation and Limitations

In this section, we investigate how agents react to a novel situation and discuss the limitations of our work.

We exemplarily increase the delay of a bottleneck edge in graph  $G_A$  from 2 to 10 at step 50 and evaluate its effect on 100 episodes of 300 steps with bandwidth limitations. Fig. 6 shows the throughput of different approaches in this scenario, combined with the stepwise mean absolute difference of node state values from the recurrent approaches. Static shortest paths (SP) ignores the delay change, stepwise SP considers it. For each learning approach, we show the results of the best model. The recurrent approaches quickly converge to small node state differences at the beginning and react to the change at step 50, although changing edge delays were never encountered during training. Before step 50, DQN performs slightly better than our approach. After the change, all three DQN models fail to adapt and display a poor throughput, while one out of three models with our approach is able to outperform stepwise shortest paths. Following research could consider dynamic graph changes during training and explore adaptivity in more detail.

These improvements in generalizability and adaptivity come at the cost of exchanging messages with all neighbor nodes at each step. Fig. 6 shows that there is comparatively little change in the node states after convergence, suggesting a reduced need for communication. While we show that a single message passing iteration per step suffices to learn graph observations for generalized routing, future work could investigate the further reduction of communication overhead. We see great potential for synergies with recent works in the area of agent-to-agent communication, where agents decide when to send messages [15, 24, 40] to selected recipients instead of broadcasting them to all other agents [26].

## 6 RELATED WORK

Reinforcement learning for graph-based environments exemplified by routing has been investigated since the introduction of Q-learning [4]. Recent works have shown to improve over previous algorithms in various domains and network conditions [22, 28].

Many of these approaches assume centralized control with a global view [1, 6, 18, 21, 35]. This not only limits their scalability, but also their reactivity. Decentralized approaches [5, 36] are more reactive, but their partial observability may degrade performance. Learning directly in the target network allows agents to specialize [35]. However, this is challenging in practice because suboptimal actions can also get stuck in local optima, requiring retraining from scratch if the graph or the network conditions change [3].

Ideally, one would pre-train agents to perform well in all graphs and network conditions and optionally fine-tune them online. Graph neural networks have shown to enable generalization in routing scenarios [8, 33], as opposed to traditional models with fixed input dimensions that specialize on individual graphs [21]. To the best of our knowledge, related works with graph neural networks are mainly restricted to centralized approaches and agent-to-agent communication [19, 29]. A noteworthy exception is the work by Geyer and Carle [10], who propose a distributed message passing scheme to learn routing in a supervised setting. With our work, we address the gap of generalizability over graphs in the context of multi-agent reinforcement learning with decentralized execution.

## 7 CONCLUSION

In this paper, we investigate the issue of generalizability in multi-agent graph environments. We propose to decouple learning graph representations and control by conceptually separating nodes and agents. Nodes iteratively learn graph representations and forward local graph observations to agents, allowing them to solve tasks in the graph. We evaluate our approach based on four graph neural network architectures across 1000 diverse graphs in a routing environment. The results indicate that recurrent graph neural networks can be trained end-to-end with reinforcement learning and sparse rewards. Graph observations do not only allow agents to generalize over different graphs, but also to adapt to changes in the graph without retraining. However, having no constraints on the resulting policies can lead to deteriorated behavior compared to agents that specialize on a single graph. This is reflected by loops in the routing environment, which we show can be alleviated with action masking.

Our contributions open up multiple avenues for future research, including the further reduction of communication overhead and the exploration of dynamically changing graphs. The effects of graph observations in different environments are also worth investigating, especially when cooperation between agents is required.

## ACKNOWLEDGMENTS

This work has been funded by the Federal Ministry of Education and Research of Germany (BMBF) through Software Campus Grant 01IS17050 (AC3Net) and the project “Open6GHub” (grant number: 16KISK014). It has been co-funded by the German Research Foundation (DFG) in the Collaborative Research Center (CRC) 1053 MAKI. The authors thank Amirkasra Amini for the valuable discussions.



## REFERENCES

- [1] Paul Almasan, José Suárez-Varela, Krzysztof Rusek, Pere Barlet-Ros, and Albert Cabellos-Aparicio. 2022. Deep reinforcement learning meets graph neural networks: Exploring a routing optimization use case. *Computer Communications* 196 (2022), 184–194.
- [2] Guillermo Bernárdez, José Suárez-Varela, Albert López, Bo Wu, Shihan Xiao, Xiangle Cheng, Pere Barlet-Ros, and Albert Cabellos-Aparicio. 2021. Is Machine Learning Ready for Traffic Engineering Optimization?. In *2021 IEEE 29th International Conference on Network Protocols (ICNP)*. 1–11.
- [3] Sai Shreyas Bhavanasi, Lorenzo Pappone, and Flavio Esposito. 2023. Dealing With Changes: Resilient Routing via Graph Neural Networks and Multi-Agent Deep Reinforcement Learning. *IEEE Transactions on Network and Service Management* 20, 3 (2023), 2283–2294.
- [4] Justin Boyan and Michael Littman. 1993. Packet Routing in Dynamically Changing Networks: A Reinforcement Learning Approach. In *Advances in Neural Information Processing Systems*, Vol. 6. Morgan-Kaufmann, 671–678.
- [5] Florian Brandherm, Julien Gedeon, Osama Abboud, and Max Mühlhäuser. 2022. BigMEC: Scalable Service Migration for Mobile Edge Computing. In *IEEE/ACM 7th Symposium on Edge Computing (SEC)*. 136–148.
- [6] Daniela M. Casas-Velasco, Oscar Mauricio Caicedo Rendon, and Nelson L. S. da Fonseca. 2021. Intelligent Routing Based on Reinforcement Learning for Software-Defined Networking. *IEEE Transactions on Network and Service Management* 18, 1 (2021), 870–881.
- [7] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. 2016. Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering. In *Advances in Neural Information Processing Systems*, Vol. 29. 3844–3852.
- [8] Miquel Ferriol-Galmés, Jordi Paillisse, José Suárez-Varela, Krzysztof Rusek, Shihan Xiao, Xiang Shi, Xiangle Cheng, Pere Barlet-Ros, and Albert Cabellos-Aparicio. 2023. RouteNet-Fermi: Network Modeling With Graph Neural Networks. *IEEE/ACM Transactions on Networking* 31, 6 (2023), 3080–3095.
- [9] Matthias Fey and Jan E. Lenssen. 2019. Fast Graph Representation Learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*.
- [10] Fabien Geyer and Georg Carle. 2018. Learning and Generating Distributed Routing Protocols Using Graph-Based Deep Learning. In *Proceedings of the 2018 Workshop on Big Data Analytics and Machine Learning for Data Communication Networks (Big-DAMA '18)*. Association for Computing Machinery, 40–45.
- [11] Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. 2017. Neural Message Passing for Quantum Chemistry. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*. PMLR, 1263–1272.
- [12] Alessio Gravina, Davide Bacciu, and Claudio Gallicchio. 2023. Anti-Symmetric DGN: a stable architecture for Deep Graph Networks. In *The 11th International Conference on Learning Representations (ICLR)*.
- [13] William L. Hamilton. 2020. Graph Representation Learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 14, 3 (2020), 1–159.
- [14] William L. Hamilton, Zhitaoying, and Jure Leskovec. 2017. Inductive Representation Learning on Large Graphs. In *Advances in Neural Information Processing Systems*, Vol. 30. 1024–1034.
- [15] Shuai Han, Mehdi Dastani, and Shihan Wang. 2023. Model-Based Sparse Communication in Multi-Agent Reinforcement Learning. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems (AAMAS '23)*. IFAAMAS, 439–447.
- [16] Eric A. Hansen, Daniel S. Bernstein, and Shlomo Zilberstein. 2004. Dynamic Programming for Partially Observable Stochastic Games. In *Proceedings of the 19th National Conference on Artificial Intelligence*. AAAI Press, 709–715.
- [17] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Computation* 9, 8 (1997), 1735–1780.
- [18] Oliver Hope and Eiko Yoneki. 2021. GDDR: GNN-based Data-Driven Routing. In *Proceedings of the 2021 IEEE 41st International Conference on Distributed Computing Systems (ICDCS)*. 517–527.
- [19] Jiechuan Jiang, Chen Dun, Tiejun Huang, and Zongqing Lu. 2020. Graph Convolutional Reinforcement Learning. In *Proceedings of the 8th International Conference on Learning Representations (ICLR)*.
- [20] Steven Kapturowski, Georg Ostrovski, John Quan, Rémi Munos, and Will Dabney. 2019. Recurrent Experience Replay in Distributed Reinforcement Learning. In *Proceedings of the 7th International Conference on Learning Representations (ICLR)*.
- [21] Gyungmin Kim, Yohan Kim, and Hyuk Lim. 2022. Deep Reinforcement Learning-Based Routing on Software-Defined Networks. *IEEE Access* 10 (2022), 18121–18133.
- [22] Tianxu Li, Kun Zhu, Nguyen Cong Luong, Dusit Niyato, Qihui Wu, Yang Zhang, and Bing Chen. 2022. Applications of Multi-Agent Reinforcement Learning in Future Internet: A Comprehensive Survey. *IEEE Communications Surveys & Tutorials* 24, 2 (2022), 1240–1279.
- [23] Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard S. Zemel. 2016. Gated Graph Sequence Neural Networks. In *Proceedings of the 4th International Conference on Learning Representations (ICLR)*.
- [24] Yen-Cheng Liu, Junjiao Tian, Nathaniel Glaser, and Zsolt Kira. 2020. When2com: Multi-Agent Perception via Communication Graph Grouping. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 4105–4114.
- [25] Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. In *7th International Conference on Learning Representations (ICLR)*.
- [26] Ziyuan Ma, Yudong Luo, and Jia Pan. 2022. Learning Selective Communication for Multi-Agent Path Finding. *IEEE Robotics and Automation Letters* 7, 2 (2022), 1455–1462.
- [27] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin A. Riedmiller, Andreas Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharmarajan Kumar, Daan Wierstra, Shane Legg, and Demis Hassabis. 2015. Human-level control through deep reinforcement learning. *Nature* 518, 7540 (2015), 529–533.
- [28] Mingshuo Nie, Dongming Chen, and Dongqi Wang. 2023. Reinforcement Learning on Graphs: A Survey. *IEEE Transactions on Emerging Topics in Computational Intelligence* 7, 4 (2023), 1065–1082.
- [29] Yaru Niu, Rohan Paleja, and Matthew Gombolay. 2021. Multi-Agent Graph-Attention Communication and Teaming. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '21)*. IFAAMAS, 964–973.
- [30] Milena Radenkovic and Vu San Ha Huynh. 2020. Cognitive Caching at the Edges for Mobile Social Community Networks: A Multi-Agent Deep Reinforcement Learning Approach. *IEEE Access* 8 (2020), 179561–179574.
- [31] Davis Rempe, Jonah Philion, Leonidas J. Guibas, Sanja Fidler, and Or Litany. 2022. Generating Useful Accident-Prone Driving Scenarios via a Learned Traffic Prior. In *Proceedings of the 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 17284–17294.
- [32] Benedek Rozemberczki, Paul Scherer, Yixuan He, George Panagopoulos, Alexander Riedel, Maria Astefanoaei, Oliver Kiss, Ferenc Beres, Guzman Lopez, Nicolas Collignon, and Rik Sarkar. 2021. PyTorch Geometric Temporal: Spatiotemporal Signal Processing with Neural Machine Learning Models. In *Proceedings of the 30th ACM International Conference on Information and Knowledge Management*. 4564–4573.
- [33] Krzysztof Rusek, José Suárez-Varela, Paul Almasan, Pere Barlet-Ros, and Albert Cabellos-Aparicio. 2020. RouteNet: Leveraging Graph Neural Networks for Network Modeling and Optimization in SDN. *IEEE Journal on Selected Areas in Communications* 38, 10 (2020), 2260–2270.
- [34] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. 2009. The Graph Neural Network Model. *IEEE Transactions on Neural Networks* 20, 1 (2009), 61–80.
- [35] Stefan Schneider, Ramin Khalili, Adnan Manzoor, Haydar Qarawlus, Rafael Schellenberg, Holger Karl, and Artur Hecker. 2021. Self-Learning Multi-Objective Service Coordination Using Deep Reinforcement Learning. *IEEE Transactions on Network and Service Management* 18, 3 (2021), 3829–3842.
- [36] Stefan Schneider, Haydar Qarawlus, and Holger Karl. 2021. Distributed Online Service Coordination Using Deep Reinforcement Learning. In *2021 IEEE 41st International Conference on Distributed Computing Systems (ICDCS)*. 539–549.
- [37] Youngjoo Seo, Michaël Defferrard, Pierre Vandergheynst, and Xavier Bresson. 2018. Structured Sequence Modeling with Graph Convolutional Recurrent Networks. In *Neural Information Processing - 25th International Conference (ICONIP) (Lecture Notes in Computer Science, Vol. 11301)*. Springer, 362–373.
- [38] Sainbayar Sukhbaatar, Arthur Szlam, and Rob Fergus. 2016. Learning Multi-agent Communication with Backpropagation. In *Advances in Neural Information Processing Systems*, Vol. 29. 2244–2252.
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems* 30. 5998–6008.
- [40] Xuefeng Wang, Xinran Li, Jiawei Shao, and Jun Zhang. 2023. AC2C: Adaptively Controlled Two-Hop Communication for Multi-Agent Reinforcement Learning. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems (AAMAS '23)*. IFAAMAS, 427–435.
- [41] Jannis Weil, Zhenghua Bao, Osama Abboud, and Tobias Meuser. 2024. Towards Generalizability of Multi-Agent Reinforcement Learning in Graphs with Recurrent Message Passing. arXiv:2402.05027 [cs.LG]
- [42] Keyulu Xu, Chengtao Li, Yonglong Tian, Tomohiro Sonobe, Ken-ichi Kawarabayashi, and Stefanie Jegelka. 2018. Representation Learning on Graphs with Jumping Knowledge Networks. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, Vol. 80. PMLR, 5453–5462.