X-GRL: An Empirical Assessment of Explainable GNN-DRL in B5G/6G Networks

Farhad Rezazadeh*, Sergio Barrachina-Muñoz*, Engin Zeydan*,

Houbing Song[‡], K.P. Subbalakshmi[§], and Josep Mangues-Bafalluy^{*}

*Centre Tecnológic de Telecomunicacions de Catalunya (CTTC), Barcelona, Spain

[‡]University of Maryland, Baltimore County (UMBC), Baltimore, USA

[§]Stevens Institute of Technology, New Jersey, USA

Contact Emails: {name.surname}@cttc.es, h.song@ieee.org, ksubbala@stevens.edu

Abstract—The rapid development of artificial intelligence (AI) techniques has triggered a revolution in beyond fifth-generation (B5G) and upcoming sixth-generation (6G) mobile networks. Despite these advances, efficient resource allocation in dynamic and complex networks remains a major challenge. This paper presents an experimental implementation of deep reinforcement learning (DRL) enhanced with graph neural networks (GNNs) on a real 5G testbed. The method addresses the explainability of GNNs by evaluating the importance of each edge in determining the model's output. The custom sampling functions feed the data into the proposed GNN-driven Monte Carlo policy gradient (REINFORCE) agent to optimize the gNodeB (gNB) radio resources according to the specific traffic demands. The demo demonstrates real-time visualization of network parameters and superior performance compared to benchmarks.

Index Terms—B5G/6G, AI/ML, XAI, GNN-DRL, Resource Allocation

I. INTRODUCTION

6G wireless communication systems herald a new era of high-speed, low-latency, reliable connectivity. They inherently support AI to pave the way for new services and use cases. Nevertheless, the opaque nature of AI can undermine trustworthiness and hinder its widespread use in critical applications. Consequently, there is a growing need for explainable AI (XAI)[1][2][3], whose primary goal is to shed light on decision-making processes. This paper investigates the feasibility of integrating Graph Convolutional Networks (GCNs) into the REINFORCE [4] algorithm in a 5G network to control the allocation of physical resource blocks (PRBs).

The main idea is that while the original state data from the network environment is not inherently structured or graphlike, our solution constructs a graph structure from this state data to enable further processing with GCNs. In our model, we conceptualize the states of the 5G network as a graph in which each state represents a node. Our main focus is to explore the graph convolution operation, a fundamental component of GCNs. In short, graph convolution is about updating node features using its characteristics and the features of neighboring nodes. This process facilitates the distribution of information across the nodes of the graph and creates a form of *communication* between them. More precisely, the operation of graph convolution takes place on the input features (x) and the edge index (edge-index) in the forward propagation process. During this process, the attributes of each node are updated



Figure 1: Infrastructure and network setup.

taking into account its neighbors, and the resulting updated node features are stored in the variable *x*. Furthermore, the importance of the edge in the graph is elucidated using the proposed explainer function that adopts an optimization-based approach. The edge mask obtained by this function provides valuable insight and explainability into the model's decision process for particular nodes in the graph.

II. TESTBED ARCHITECTURE

The testbed is shown in Fig. 1 and consists of specific hardware and software components to realize the 5G infrastructure and management-related functionalities. For the implementation of the gNBs and the 5G core, the testbed uses a Callbox Ultimate and a Callbox Mini, respectively. To emulate user equipment (UE) interactions with the 5G network, the testbed incorporates a Simbox. These Amarisoft components were selected for their ability to accurately replicate the behavior of 5G elements in a controlled environment. The physical setup of RAN with the Amarisoft devices is shown in Fig. 2.

On top of this 5G infrastructure, the testbed deploys a cloud-native platform. This Kubernetes-based platform hosts



Figure 2: 5G RAN components.

three main management-related workloads. First, the API client processes requests to the 5G network elements. Second, the monitoring system continuously collects relevant metrics from all components of the system to help with real-time performance analysis and system evaluation. This monitoring system is implemented as a collection of pods that contain a Kafka bus along with custom sampling functions that span multiple domains. And finally, the decision engine performs actions at the RAN level. As for the data plane, the testbed uses MGEN ¹[5] to generate downlink traffic patterns of interest.



Figure 3: Traffic patterns generated throughout the demo.

III. OPERATIONAL PHASES

As shown in Fig. 3, we examine two different traffic patterns, which are then forwarded to Simbox as downlink traffic. The pattern A has a stochastic character. In this case, the packets arrive according to a Poisson process, where the arrival time between packets corresponds to an exponential distribution. This means that packets can arrive at any time, regardless of when the previous packet was received. On the other hand, the traffic pattern B adheres to a periodic pattern and transmits data packets at regular intervals. This regularity facilitates learning because of the predictable data rate.

The decision engine, developed in Python, uses the OpenAI Gym²[6] toolkit to interface with DRL agents. Fig. 4 illustrates the average reward per episode for five different algorithms with a single UE. The reward function indicates the effectiveness of decision making, with higher rewards indicating a more optimal allocation method. Our proposed GNN-REINFORCE algorithm ensures superior learning generalization and robust performance

²https://gymnasium.farama.org/



Figure 4: Performance comparison of various DRL approaches for radio resource allocation in the training phase.

compared to other DRL approaches. At the beginning of the training, the agent starts exploring the action space, i.e. the PRBs chunks. This concept of exploration is ingrained through the probabilistic aspect of action selection, a fundamental feature of policy gradient techniques. The agent then tries to find an optimal balance between the learned decision policies and the network states, indicating the exploitation phase. The inclusion of a graph convolution operation in the GNN-REINFORCE algorithm significantly improves the performance of the agent compared to the standard REINFORCE algorithm.

Fig. 5 gives an insight into the explainability of the GNN-REINFORCE algorithm by highlighting the importance of edge connections. This analysis explores the dynamic evolution of edge importance values over time within the GNN. It sheds light on the learning process of the network and how it gradually prioritizes edges between nodes. During the initial training phase, we observe that all edges have an equal importance value of one. This indicates the algorithm is starting to learn the relationships between the nodes. We notice a variation in the edge importance values as training progresses to the midtraining phase. The weights are starting to be adjusted on the feedback received from the network environment. In the posttraining or inference phase, a transformation is evident. Only one edge is marked as important, with a non-zero value, while the importance values of all other edges have diminished to zero. This evolution signifies the key edge (or relationship) that is most influential in making decisions or predictions.

Fig. 6-(a) and (b) show the network performance of the GNN-REINFORCE algorithm under Traffic-A and Traffic-B, respectively. Each plot is divided into two main sections showing the cumulative distribution function (CDF) of the gap in the allocation of radio resources and the agent's performance in the allocation of PRBs. In 5G network management, the radio resource allocation gap is typically defined as the discrepancy between the expected or desired allocation of a radio resource and the actual allocation, which directly affects the overall performance of the network.

The CDF provides a statistical assessment of the efficiency of resource allocation. Zero gap values imply ideal allocation, indicating more effective use of resources. The performance of the GNN-REINFORCE agent demonstrates a remarkable ability to allocate PRBs. It allocates resources according to traffic demand with a remarkable accuracy of more than 90%.

In Fig. 7, we evaluated the robustness of our GNN-REINFORCE model by introducing Gaussian noise into the

¹https://github.com/USNavalResearchLaboratory/mgen



(a) Early training stage

(b) Mid-Training stage

(c) Post-Training analysis

Figure 5: Explainability analysis of the GNN-REINFORCE algorithm.



(b) Under Traffic-B.

Figure 6: Analysis of network performance under Traffic A and B. (Up) The CDF of radio resource allocation gap, (Down) Performance evaluation in terms of PRB allocation.



Figure 7: Robustness of the model in inference phase.

input state values, with noise values ranging from 0 to 0.1. We established these disturbance levels by creating 20 numbers evenly distributed in this range, each representing the standard deviation of the Gaussian noise. We then exposed the model to each disturbance level and logged its performance. The objective of this evaluation was to investigate the ability of the model to maintain its performance at various noise disturbances. As shown in Fig. 7, the model showed satisfactory performance in the inference phase compared to the average reward at convergence. Fig. 8 shows the inference phase and the main metrics monitored and observed on Grafana dashboard.



Figure 8: Visualization of metrics of the inference phase with Grafana.

IV. CONCLUSION

The synergy of DRL, GNNs, and XAI offers a compelling path towards enhanced efficiency, transparency, and reliability in mobile network management. Future work will extend this approach to larger and more complex network scenarios and explore potential applications in other areas of network operations and management.

ACKNOWLEDGMENT

This work was partially funded by MCIN/AEI/ 10.13039/501100011033 grant PID2021-1264310B-I00 (ANEMONE), Spanish MINECO grant TSI-063000-2021-54 (6G-DAWN) and grant TSI-063000-2021-56 (6G-BLUR), Generalitat de Catalunya grant 2021 SGR 00770 (6GE2E),

REFERENCES

- [1] W. Tan *et al.*, "NoiseCAM: Explainable AI for the Boundary Between Noise and Adversarial Attacks," *arXiv:2303.06151*, 2023.
- [2] F. Rezazadeh, H. Chergui, L. Alonso, and C. Verikoukis, "SliceOps: Explainable MLOps for Streamlined Automation-Native 6G Networks," arXiv:2307.01658, 2023.
- [3] M. Chen, N. Wang, and K. P. Subbalakshmi, "Explainable Rumor Detection using Inter and Intra-feature Attention Networks," *TrueFact KDD Workshop*, 2020.
- [4] J. Zhang et al., "Sample Efficient Reinforcement Learning with REIN-FORCE," The Thirty-Fifth AAAI Conference on AI, 2021.
- [5] F. Rezazadeh, L. Zanzi, F. Devoti, S. Barrachina-Munoz, E. Zeydan, X. Costa-Pérez, and J. Mangues-Bafalluy, "A Multi-Agent Deep Reinforcement Learning Approach for RAN Resource Allocation in O-RAN," *IEEE INFOCOM*, 2023.
- [6] F. Rezazadeh, L. Zanzi, F. Devoti, H. Chergui, X. Costa-Pérez, and C. Verikoukis, "On the Specialization of FDRL Agents for Scalable and Distributed 6G RAN Slicing Orchestration," *IEEE Transactions on Vehicular Technology*, vol. 72, no. 3, pp. 3473–3487, 2022.