Global-Local Attention vs Graph Neural Networks in the Reinforcement Learning Approach for the Dynamic Berth Allocation Problem

Bridging the Optimality Gap in Dynamic Berth Allocation Problem via Global-Local Attention

Vlad-Ioan Anica-Popa[1]
Supervisor(s): Neil Yorke-Smith[1]**, Carlos March Moya**[1]
[1]EEMCS, Delft University of Technology, The Netherlands

A Thesis Submitted to EEMCS Faculty Delft University of Technology,
in Partial Fulfilment of the Requirements
for the Bachelor of Computer Science and Engineering
21 June 2026

Name of the student: Vlad-Ioan Anica-Popa
Final project course: CSE3000 Research Project
Thesis committee: Neil Yorke-Smith, Carlos March Moya, Wendelin Böhmer

Global-Local Attention vs Graph Neural Networks in the Reinforcement Learning Approach for the Dynamic Berth Allocation Problem

Author: Vlad-Ioan Anica-Popa

Supervisors: Neil Yorke-Smith, Carlos March Moya

Abstract

The Dynamic Berth Allocation Problem (DBAP) is an optimization problem in maritime logistics that seeks to minimize vessel delays and improve terminal efficiency through effective berth scheduling. This paper investigates how replacing the Graph Neural Network (GNN) encoder in a reinforcement learning approach to the DBAP with a Global-Local Attention mechanism affects the optimality gap. This new encoder architecture is designed to capture both local vessel-berth interactions and global scheduling information. The performance of the proposed architecture is evaluated against the baseline GNN across 2,700 benchmark instances spanning various terminal sizes and traffic congestion levels. Empirical results demonstrate that the Global-Local Attention variant yields a consistent improvement in schedule quality, reducing normalized operational costs by an average of 2.8%. While this modification provides notable optimization gains, particularly in lower congestion scenarios, it introduces an 81.4% increase in trained parameters, presenting a distinct trade-off between scheduling optimality and computational efficiency.

1 Introduction

In the era of economic globalization and trade liberalization, maritime transportation serves as the fundamental backbone of the global supply chain, positioning seaports as key nodes of international commerce. In 2025, over 80% of shipped goods were transported via sea transportation [1]. As a result, the growing number of large cargo vessels has significantly increased traffic at terminals around the world, revealing major capacity limitations in port infrastructure [2]. These challenges are made worse by ongoing operational disruptions, such as changing vessel routes and less reliable schedules. Empirical results show that between December 2023 and December 2024 sea vessel waiting times climbed from 5.22 hours to 6.39 hours, an increase of 22.5% [3]. In this context, minimizing vessels' turnaround time is crucial for reducing maritime congestion, accelerating overall cargo transit, and stabilizing supply chain reliability. Since vessels spend a significant portion of their time waiting for berth availability, improving berth allocation decisions can directly impact these metrics.

To evaluate terminal performance and formulate optimization objectives, port authorities rely heavily on vessel turnaround time as a definitive metric of operational efficiency. As conceptualized by [4], this key performance indicator encompasses the entire temporal window from a vessel's entry into port limits to its subsequent departure. Structurally, turnaround time is disaggregated into three primary operational segments: (i) waiting time at anchorage prior to docking, (ii) maneuvering latencies during the berthing and unberthing phases, and (iii) active berth time, during which service operations such as loading and unloading occur.

From a reinforcement learning perspective, the sum of the mentioned times coupled with a vessel's priority contribute to the function the agent aims to minimize, thus directly influencing the agent's rewards. Consequently, minimizing this component (waiting time combined with a vessel's priority) is essential to maximizing a terminal's throughput and reducing overall port stay latencies. Such operational efficiencies yield broader sustainability benefits: by accelerating port turnaround and reducing idle stay times, vessel operators gain the operational flexibility to implement speed reduction strategies at sea which can decrease energy consumption near ports by up to 25.4% [5]. In addition, reducing vessel delays helps address important environmental issues, such as worsening air quality in nearby port cities, requiring terminal operators to balance efficient operations with emissions reduction goals. From an economic perspective, shipping companies depend on smooth port operations because unnecessary waiting increases total travel time, raises daily operating and crew costs, and lowers how efficiently vessels are used, leading to both higher costs and greater environmental impact over time [6].

To improve operations, increase efficiency, and provide better service, port authorities usually follow one of two main strategies: expanding physical infrastructure through major investments or using advanced scheduling models to make better use of existing berths, anchorages, and waterways [6]. Compared to the high costs and long timelines of infrastructure expansion, optimization algorithms offer a more sustainable, efficient, and practical way to improve port performance.

This paper focuses on improving the allocation of quay space to incoming vessels, determining where vessels should berth to perform loading and unloading operations at the terminal. This problem is called the Berth Allocation Problem. As described in Bierwirth and Meisel [7], the BAP involves determining the berth and the time a vessel will be serviced at. It considers berth-specific handling times and other operational constraints. The variation of the problem that we are trying to solve is the Dynamic Berth Allocation Problem (DBAP), where vessels arrivals occur progressively throughout the planning horizon, rather than all being present at its start.

In this work, we address the DBAP using reinforcement learning. Our approach builds upon the framework proposed in March Moya et al. [8], which formulates the DBAP as a Markov Decision Process (MDP), enabling the problem to be solved through sequential decision-making techniques. To represent the state of the environment, the authors employ a Graph Neural Network (GNN)-based architecture that generates embeddings for both vessels and berths, allowing the model to capture the relationships between them. The effectiveness of the proposed approach is demonstrated through an empirical evaluation on 2,700 randomly generated problem instances, where the reinforcement learning model consistently outperforms traditional dispatching heuristics in terms of solution quality. These results highlight the potential of learning-based methods for solving complex berth allocation problems and provide the foundation upon which this work builds.

While the work of March Moya et al. [8] demonstrates that reinforcement learning combined with GNN-based embeddings can outperform traditional dispatching heuristics, it remains unclear whether graph neural networks are the most effective representation learning mechanism for the DBAP. Advances in Global-Local Attention architectures have shown improved capability in capturing both local interactions and global contextual information in Beltagy et al. [9], however, their applicability to dynamic berth allocation has not yet been investigated.

The research question this paper aims to answer is whether the Global-Local Attention mechanism reduces the optimality gap in the Dynamic BAP compared to the GNN approach. Therefore, the primary contribution of this work is the replacement of the existing Graph Neural Network (GNN)-based architecture used for generating vessel and berth embeddings with a Global-Local Attention architecture, as proposed in Beltagy et al. [9]. The motivation behind this modification is to investigate whether the Global-Local Attention mechanism can capture both local interactions and global structural information more effectively than the GNN-based approach, thereby producing richer and more informative representations of vessels and berths.

The main contributions of this work are:

- The integration of a Global-Local Attention architecture into an existing MDP for the DBAP.
- An empirical comparison between Global-Local Attention and GNN-based vessel and berth embeddings.
- An implementation of the Global-Local Attention mechanism [10].

The remainder of this paper is organized as follows. Section 2 reviews related work and provides background on the DBAP and reinforcement learning approaches. Section 3 describes the baseline GNN architecture and the proposed Global-Local Attention model. Section 4 presents the experimental setup and evaluation methodology. Section 5 discusses the responsible research aspect. Section 6 discusses possible reasons for the found results and limitations. Finally, Section 7 concludes the paper and outlines directions for future research.

2 Background

According to Lv et al. [11], UNCTAD [12] port resilience refers to the ability of a port to maintain an acceptable level of operational performance during disruptions and to recover efficiently once normal conditions are restored. Such disruptions may arise from adverse weather conditions, equipment failures, labour shortages, unexpected surges in vessel arrivals, or broader supply chain disturbances. In these situations, effective berth allocation plays a vital role in mitigating congestion and preventing delays from propagating throughout the port. By efficiently assigning vessels to available berths and adapting to changing operational conditions, a berth allocation system can help maintain throughput and minimize the impact of disruptions. Consequently, resilient berth allocation strategies contribute not only to the stability of individual ports but also to the reliability of global supply chains and international trade networks.

The Berth Allocation Problem (BAP) is a foundational challenge in maritime logistics, tasking port operators with the optimal assignment of arriving vessels to discrete or continuous berthing spaces over a temporal horizon. Given its inherently NP-hard nature [13], previous literature has approached the BAP through a diverse spectrum of methodologies, broadly categorized into exact mathematical programming and heuristic optimization frameworks.

2.1 Related Work solving the BAP

Imai et al. [14] applied Lagrangian relaxation, however, this solution does not scale well when increasing the size of the problem's parameters, which means it is not applicable for large-scale ports. Another approach is presented in [15], where a branch-and-bound algorithm (Sedimentation Algorithm with and Estimation and Rearrangement Heuristic) was used. By leveraging pruning mechanisms, their method tightened bounds more effectively than standard exact solvers. However, despite these algorithmic improvements, this methodology encounters the same combinatorial explosion characteristic of exact mathematical programming. Lin et al. [16] introduced an iterated greedy heuristic targeting the minimization of total service time in discrete dynamic berth allocation problems. Their framework proved highly competitive against traditional metaheuristics, consistently identifying optimal or best-known solutions across complex, large-scale test instances. Türkoğulları et al. [17] proposes an integrated optimization framework that simultaneously determines vessel berthing locations and times, assigns quay cranes to vessels over time, and schedules crane operations while accounting for crane relocation costs, thereby improving the realism of berth allocation models and remaining applicable to practical problem instances.

There were other approaches using metaheuristic algorithms. One of them is [18] which made use of Ant Colony Optimization (ACO). Another one is [19] which leveraged a Genetic Algorithm (GA) to assign vessels to berths, while also trying to optimize for certain vessel preferences (e.g. daytime preference). Hsu [20] developed a non-linear Mixed Integer Programming model that integrated the discrete berth allocation problem with dynamic quay crane assignments. By introducing a Hybrid Particle Swarm Optimization (HPSO) paired with an event-based heuristic, the study successfully transitions away from rigid, time-invariant scheduling to a flexible, variable-in-time crane assignment framework. Experimental results demonstrated that this swarm-based framework statistically outperformed traditional Genetic Algorithms in reducing vessel delays and operational costs. de Oliveira et al. [21] studied the DDAP and proposed a hybrid metaheuristic combining Clustering Search and Simulated Annealing. Their approach modeled berth allocation as a vehicle routing problem with time windows and uses clustering mechanisms to identify and intensify the search within promising regions of the solution space. Computational results on benchmark instances demonstrate that the method achieved state-of-the-art solution quality at the time of publication while maintaining relatively low computational requirements. Bacalhau et al. [22] proposed two hybrid metaheuristics, GASSR and MASSR, for the Dynamic Berth Allocation Problem, combining genetic algorithms with an approximate dynamic programming local search procedure. Their model extended traditional DBAP formulations by incorporating practical port-specific constraints such as berth cargo preferences, safety requirements, and heterogeneous berth characteristics derived from real-world operations. Computational results on benchmark instances based on data from the Port of Paranaguá showed that the proposed hybrid methods outperform standard genetic algorithms, particularly for large-scale and tightly constrained scheduling scenarios. Lin and Ting [23] developed two simulated annealing-based metaheuristics for the Dynamic Berth Allocation Problem, including a variant augmented with a restart strategy to improve diversification during the search process. Their approach employs a permutation-based berth scheduling representation together with swap and insertion neighborhood moves, enabling efficient exploration of large solution spaces. Computational experiments produced numerous new best-known solutions on benchmark continuous berth allocation instances. Lalla-Ruiz et al. [24] proposed a hybrid metaheuristic for the Dynamic Berth Allocation Problem that integrates Tabu Search with Path Relinking, leveraging an elite set of high-quality solutions to intensify and diversify the search process. The method extends earlier Tabu Search approaches through the use of multiple neighborhood structures, including vessel reallocation and swap moves, and generates new starting solutions via path relinking between elite and randomly generated solutions. Computational results showed that the hybrid approach significantly outperformed previous tabu search methods and produced near-optimal solutions even for large-scale instances where exact optimization methods become computationally infeasible.

2.2 Reinforcement Learning

Reinforcement Learning (RL) is a branch of machine learning in which an agent learns to make sequential decisions by interacting with an environment and receiving feedback in the form of rewards. At each decision step, the agent observes the current state of the environment, selects an action according to a policy, and transitions to a new state while obtaining a reward signal. The objective of the agent is to maximize the cumulative reward over time, thereby learning behaviors that optimize long-term outcomes rather than immediate gains. RL problems are commonly formulated as Markov Decision Processes, which define the environment through states, actions, transition dynamics, and reward functions. This formulation makes reinforcement learning particularly suitable for complex optimization and scheduling problems, where decisions influence future system states and must be adapted dynamically over time. Recent advances combining reinforcement learning with deep neural networks have further expanded its applicability to large-scale decision-making problems and combinatorial optimization tasks.

2.3 Global-Local Attention

While GNNs effectively capture information from neighboring nodes, information originating from distant parts of the graph requires multiple message-passing layers to propagate. This can make it difficult to model global dependencies within large scheduling instances. Recent attention-based architectures address this limitation by allowing selected nodes to directly exchange information with a much larger portion of the graph.

The Global-Local Attention mechanism used in this work is inspired by the sparse attention strategy introduced in Beltagy et al. [9]. Instead of performing full self-attention between all pairs of nodes, which would be computationally expensive, only a subset of nodes is designated as global nodes. These nodes influence the embeddings of all the other nodes in the graph and are themselves influenced by all the remaining nodes. The rest of the nodes continue to interact primarily through local neighborhood connections.

This approach seeks to combine the strengths of local message passing and global information exchange. Local interactions preserve the structural information captured by the original GNN, while global attention enables information from distant graph regions to be incorporated more directly.

3 Replacing GNNs with Global-Local Attention

3.1 Baseline GNN Architecture

The reinforcement learning framework proposed in March Moya et al. [8] models the DBAP as a Markov Decision Process (MDP), as shown in Figure 1. At every decision step, the current state of the terminal is represented as a bipartite graph consisting of two node types:

vessels and berths. An edge is present between a vessel and a berth whenever the berth is capable of servicing that vessel.

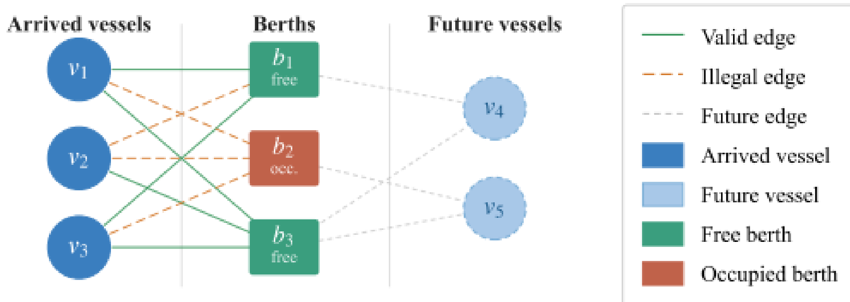

Figure 1: Representing a unified bipartite state graph from March Moya et al. [8]

To obtain a numerical representation of the current state, the framework employs a graph neural network. The GNN iteratively exchanges information between connected nodes through message passing. During each layer, vessel nodes aggregate information from neighboring berth nodes, while berth nodes aggregate information from neighboring vessel nodes. As a result, the final node embeddings contain information about both local node attributes and the structure of the surrounding graph.

The resulting vessel and berth embeddings are subsequently processed by a policy network which determines the next action. Possible actions include assigning a vessel to a berth or postponing the decision until the next event. In the DBAP environment, time advances in an event-driven manner rather than through fixed-length timesteps. An event corresponds to a significant change in the system state, such as the arrival of a new vessel or the completion of service at a berth. When the agent chooses to wait, the simulation progresses to the next such event, after which a new decision can be made based on the updated state of the terminal. The quality of the policy is optimized using reinforcement learning, where rewards are based on the operational objectives of the berth allocation problem.

3.2 Selection of Global Nodes for Global-Local Attention

A crucial design choice in the Global-Local Attention architecture is determining which nodes should receive global attention. In the context of berth allocation, we hypothesize that certain berths have a disproportionate impact on overall scheduling performance. Consequently, the global nodes are selected from the set of berth nodes rather than vessel nodes. More specifically, we choose the first $\sqrt{N}$ berths with the lowest average handling times, where N denotes the total number of berth nodes in the graph.

The motivation behind this choice is that highly efficient berths are expected to process a larger number of vessels throughout the planning time horizon. Decisions involving these berths may therefore have a greater influence on total waiting times and overall schedule quality. By allowing these nodes to communicate globally, the model may be able to better coordinate berth assignments across the entire graph.

3.3 Proposed Architecture

The proposed architecture retains the overall reinforcement learning framework and policy network of the original approach. The only modification is the replacement of the pure GNN encoder with a hybrid Global-Local Attention encoder, as shown in Figure 2, where the node circled in green represents a high performing berth that influences the entire graph. Initially, vessel and berth features are embedded into a shared latent space. Message passing is then performed using the Global-Local Attention mechanism. Local interactions occur along the edges of the bipartite graph, while the selected global berth nodes additionally exchange information with all nodes in the graph. After several encoding layers, the resulting node embeddings are forwarded to the same policy network used by the baseline method.

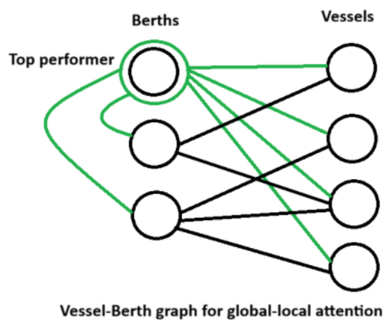

Figure 2: Bipartite graph representing state.

By keeping all other components unchanged, any observed performance differences can be attributed directly to the change in representation learning. This allows for a fair comparison between the original GNN encoder and the proposed Global-Local Attention encoder.

If this hypothesis holds, the reinforcement learning agent should produce berth allocation schedules with a smaller optimality gap than those generated using the original GNN-based encoder. The experiments presented in the following section evaluate this hypothesis empirically.

4 Experimental Setup and Results

This section presents the empirical evaluation comparing the proposed Global-Local Attention model against the baseline (GNN) approach. The Global-Local Attention variant yields an average performance gain of approximately 2.8% across the 2,700 unique test instances. A detailed breakdown of performance metrics, statistical significance tests, and an analysis of model scalability across varying terminal configurations are provided below.

4.1 Experimental Setup

Vessel arrivals: Following the methodology established by [25], vessel inter-arrival times are modeled using an exponential distribution with a rate of $\lambda = \frac{1}{I}$. The parameter I is defined as:

$$I = C \cdot \frac{\sum_{v,b} h_{v,b}}{B \cdot V^2}$$

Here, $C > 0$ is a congestion control coefficient, where smaller values of C correspond to higher levels of traffic congestion.Vessel attributes: handling times for each vessel are uniformly sampled from $U[10, 200]$, and vessel priorities are drawn uniformly from the discrete set $\{1, \ldots, 10\}$.Training and validation configurations: Training batches alternate between small configurations (10 vessels, 3 berths) and medium configurations (30 vessels, 5 berths) across congestion levels $C \in \{0.5, 1.0, 2.0\}$. This yields 12 instances per batch, with each instance evaluated over $E_{inst} = 5$ episodes. To ensure consistent tracking during training, the validation set consists of six fixed instances spanning the same size and congestion configurations.Test benchmarks: To rigorously evaluate scalability and generalization, the test set encompasses all 27 combinations of total vessels $V \in \{80, 100, 120\}$, available berths $B \in \{5, 10, 15\}$, and congestion levels $C \in \{0.5, 1.0, 2.0\}$. With 100 independent randomized replicates per combination, the final test suite comprises 2,700 unique instances.

4.2 Results

When compared to the baseline Graph Neural Network approach, the Global-Local Attention variant demonstrates a consistent, improvement in solution quality, achieving an average performance gain of approximately 2.8% across the test set. As illustrated in Figure 3 and detailed in Table 1, the distribution of results indicates that while both approaches achieve comparable performance on a majority of instances, the Global-Local Attention model successfully identifies significantly better schedules than the GNN model in specific cases, with some outliers showing performance improvements exceeding 10%.

A deeper analysis of the environmental parameters reveals distinct scenarios where the proposed architecture excels. While the first graph from Figure 4 shows that the number of vessels does not change the improvement in performance when compared to the GNN agent, the second graph demonstrates how the model scales exceptionally well with terminal size, peaking at an average improvement of 5.31% when 15 berths are available. Furthermore, the last graph of Figure 4 demonstrates that the Global-Local Attention model achieves its most significant performance gains in lower congestion scenarios ($C = 2.0$), yielding an average improvement of 4.32%. Conversely, improvements are more muted under high congestion ($C = 0.5$).

These performance trends can be directly attributed to the architecture of the Global-Local Attention mechanism. Unlike standard GNNs that rely purely on local message passing, the proposed model explicitly designates a subset of nodes as "global nodes" capable of attending to the entire graph. Specifically, the framework selects the $\sqrt{N}$ berths with the lowest average handling times to act as these global communication hubs.

In scenarios characterized by higher berth availability (e.g., 15 berths) and lower congestion, vessels have a wider array of feasible scheduling options and greater flexibility in avoiding queues. Under these conditions, the global attention mechanism effectively broadcasts the availability and state of the terminal's most performative berths to all vessels simultaneously. This structural advantage allows the reinforcement learning agent to encode richer, more actionable state representations, leading to optimal vessel-to-berth routing decisions that a purely localized GNN might miss.

To determine whether the observed performance difference between the two architectures was statistically meaningful rather than a result of random variation, a paired statistical significance test was conducted on the evaluation results. Specifically, a paired t-test was performed to compare the normalized costs produced by the Global-local Attention variant and the baseline GNN variant across all benchmark instances. The test yielded a p-value of 0.002, which is below the chosen significance threshold of 0.05. Therefore, the null hypothesis that both models achieve equivalent performance was rejected, indicating that the Global-Local Attention variant achieved a statistically significant improvement over the GNN-based approach.

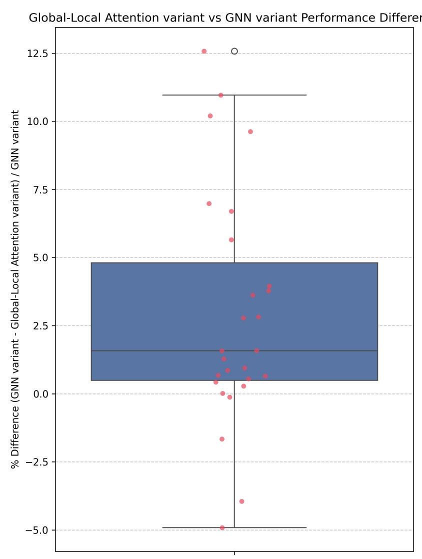

Figure 3: Boxplot of differences in performance between the two models in percentages, where the baseline is the GNN approach

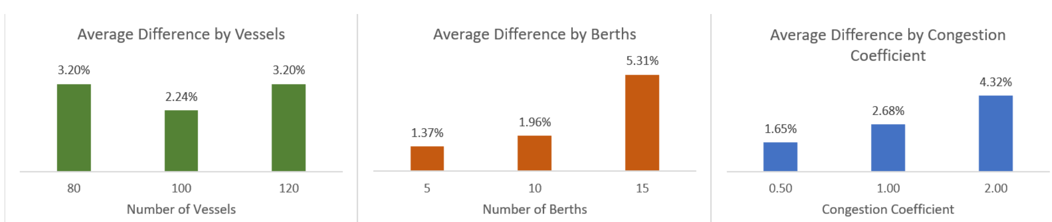

Figure 4: Aggregate percentage improvements across different instance-generation parameters, grouped respectively by the number of vessels, the number of berths, and the congestion coefficient

5 Responsible Research

The primary socio-economic consideration of this work centers on port automation and labor dynamics. Rather than viewing the reinforcement learning framework as an autonomous

Table 1: Percentage difference of the Global-Local Attention variant (Model B) over the GNN variant (Model A). Positive values indicate the Global-Local Attention variant achieved a lower (better) normalized cost.

Vessels	Berths	Congestion coefficient	GNN variant	Global-Local Attention variant	Difference (%)
80	5	0.5	896.48	892.66	0.43
80	5	1	1163.95	1156.03	0.68
80	5	2	918.55	883.82	3.78
80	10	0.5	343.65	357.21	-3.95
80	10	1	371.81	345.88	6.98
80	10	2	269.06	282.29	-4.92
80	15	0.5	228.90	200.12	12.57
80	15	1	232.10	209.76	9.62
80	15	2	185.45	178.72	3.63
100	5	0.5	1505.61	1497.28	0.55
100	5	1	1202.57	1222.58	-1.66
100	5	2	1297.45	1276.85	1.59
100	10	0.5	551.26	546.54	0.86
100	10	1	427.27	427.81	-0.13
100	10	2	387.91	345.40	10.96
100	15	0.5	303.41	294.97	2.78
100	15	1	263.94	253.51	3.95
100	15	2	178.82	176.53	1.28
120	5	0.5	1734.61	1734.40	0.01
120	5	1	1592.68	1588.25	0.28
120	5	2	1281.59	1195.82	6.69
120	10	0.5	523.84	520.44	0.65
120	10	1	509.08	501.07	1.57
120	10	2	375.53	354.30	5.65
120	15	0.5	247.57	245.23	0.94
120	15	1	243.56	236.68	2.82
120	15	2	221.01	198.46	10.20

replacement for human port authorities, the proposed model is designed strictly as a decision-support tool. Because the simulation operates on an event-driven basis, its outputs provide strategic scheduling recommendations that human operators can audit, override, or adapt in real time to accommodate unmodeled maritime safety constraints or sudden labor disruptions. Furthermore, there is an environmental imperative tied to this work. By reducing the optimality gap, we effectively minimize vessel waiting times and idle engine hours, which directly contributes to a reduction in carbon emissions at the port. However, we also acknowledge the environmental cost of computation; the training of deep reinforcement learning models requires significant energy.

To ensure that the results of this study can be independently verified and built upon by the research community, we adhere to the following reproducibility standards:

- Algorithmic Transparency: Detailed architectures for both the baseline Graph Neural Network and the proposed Global-Local Attention mechanism are provided, including layer dimensions, activation functions, and attention head configurations.
- Data Availability: The synthetic datasets used to simulate vessel arrivals and berth constraints are generated via a documented process, and the source code for the simulation environment and training pipeline will be made available via a public repository (GitHub)[10].

By using reinforcement learning as a decision-support tool instead of a fully automatic system, this approach aims to improve terminal performance while following Human-In-The-Loop (HITL) principle. These ideas help ensure that better scheduling also supports environmental goals, reliable results, and practical use in port operations.

6 Discussion and Limitations

This study evaluates the integration of a Global-Local Attention encoder into a reinforcement learning framework for the Dynamic Berth Allocation Problem (DBAP), comparing it against a baseline Graph Neural Network (GNN) across 2,700 benchmark instances spanning various congestion levels and terminal sizes. The empirical evaluation concluded that the Global-Local Attention mechanism provides a consistent advantage, reducing the optimality gap by an average of 2.8% across the test suite.

These findings contribute to the evolving landscape of DBAP optimization. The reinforcement learning approach equipped with GNNs [8] demonstrates that autonomous agents could outperform traditional heuristics. By implementing a sparse attention strategy inspired by recent natural language processing architectures[9], this work demonstrates that selectively broadcasting information from high-impact nodes can further enhance representation learning in complex maritime logistics, allowing the agent to anticipate global scheduling dynamics.

The performance distribution reveals a strong interaction between the terminal's operational state (congestion and berth count) and the efficacy of the global attention mechanism. The model demonstrates the highest performance gains in scenarios characterized by lower congestion ($C = 2.0$) and a larger number of available berths. A theoretical explanation for this outcome is the spacial and temporal flexibility afforded to the reinforcement learning agent in the given context. When congestion is low and berths are plentiful, vessels have multiple viable routing options. By utilizing the $\sqrt{N}$ fastest berths as globally connected nodes, the attention mechanism effectively broadcasts the status of the terminal's most efficient resources to the entire graph. This allows the agent to strategically prioritize long-term efficiency, routing ships to optimal berths without the immediate pressure of overwhelming queues.

Conversely, the Global-Local Attention model's advantage diminishes in highly congested scenarios ($C = 0.5$). In heavily backlogged terminals, the scheduling problem becomes fundamentally constrained by a lack of capacity rather than a lack of routing optimization; an agent must utilize any available berth simply to clear the queue. In such rigidly constrained environments, the global communication from the fastest berths becomes saturated and less actionable. The problem reverts to relying heavily on immediate, localized vessel-berth compatibilities, a task that standard GNN message passing is already equipped to handle.

Furthermore, the strategy of selecting the fastest berths as global hubs assumes that handling time is the universal driver of schedule quality. While this holds true when capacity is abundant, future iterations of this model might benefit from dynamically learning which nodes should act as global hubs based on the real-time structural centrality of the terminal's state, rather than relying on a static, predefined selection.

A dimension in evaluating the viability of the Global-Local Attention model for real-time maritime logistics is the trade-off between scheduling optimization and computational overhead. While the proposed architecture yields a demonstrable performance improvement, reducing normalized operational costs by an average of 2.8%, it does so at the expense of significantly increased structural complexity. Specifically, the baseline GNN variant utilizes a configuration of 170,820 trained parameters. In contrast, the integration of global communication hubs increases the trained parameter count to 309,956, representing an approximate 81.4% increase in the model's total learnable capacity. While the parameter count is low for modern standards, it is worth noting that the improvement in performance of 2.8% did come at an increase of 81.4% increase in parameter count. A higher parameter count implies higher training times for the agent and higher computational costs. This aspect should be considered when choosing between the two solutions.

A limitation of this study is the synthetic generation of the benchmark instances. In the current experimental setup, handling times are sampled independently for every vessel-berth pair. Consequently, the berths identified as the most efficient may not exhibit substantially different characteristics from the remaining berths. Since the proposed method relies on selecting the fastest berths as global nodes, the benefits of allowing these nodes to exchange information globally may be artificially limited by this data homogeneity. If no berth consistently plays a dominant role in the scheduling process, prioritizing certain berths through global attention is unlikely to provide a significant advantage or fully justify the 80% increase in trained parameters.

Another limitation is that only a single strategy for selecting global nodes was evaluated. The choice of assigning global attention to the $\sqrt{N}$ fastest berths was motivated by the assumption that these berths exert the greatest influence on the overall schedule quality, however, this assumption was not empirically validated. Alternative selection mechanisms may yield different outcomes, including assigning global attention based on vessel priorities, graph centrality measures, berth utilization patterns, or allowing the model to learn the global nodes dynamically during training. Exploring these alternatives could provide further insight into whether the observed performance gains originate from the Global-Local

Attention mechanism itself or from the particular node-selection strategy employed in this work.

7 Conclusion and Future Works

This paper investigates whether replacing the Graph Neural Network (GNN) encoder used in a reinforcement learning framework for the Dynamic Berth Allocation Problem (DBAP) with a Global-Local Attention architecture could reduce the optimality gap of generated berth allocation schedules. The proposed approach introduces globally connected berth nodes inspired by sparse attention mechanisms, with the objective of combining local vessel-berth interactions with global scheduling information.

Empirical evaluations across 2,700 benchmark instances demonstrate that the Global-Local Attention mechanism provides a consistent advantage, reducing the optimality gap by an average of 2.8% compared to the baseline GNN approach. The architecture excels particularly in scenarios characterized by lower congestion and a larger number of available berths. In these environments, utilizing the fastest berths as global communication hubs effectively broadcasts the status of the terminal's most efficient resources to the entire graph, allowing the agent to prioritize long-term efficiency. However, this improvement in schedule quality involves a computational trade-off: the attention-based model increases the structural complexity, requiring an 81.4% increase in trained parameters to achieve these gains.

To address the limitations and further validate the architecture, several promising directions for future research are identified:

- Enhanced structural heterogeneity: future work should evaluate the approach on benchmark instances with stronger structural heterogeneity. One promising direction would be to assign each berth an inherent efficiency level and each vessel a difficulty level, deriving handling times from these attributes. Such a generation process would create clearer distinctions between important and less important berths, potentially allowing the benefits of global attention to emerge more strongly.
- The method used to select global nodes may influence the results. The choice of the $\sqrt{N}$ fastest berths is motivated by the assumption that these berths would have the greatest impact on schedule quality. However, this assumption is not explicitly validated. Other node-selection strategies may better capture the most influential components of the scheduling process. For example, global attention could be assigned to vessels with high priority values, to nodes with high graph centrality, or even learned dynamically by the model itself rather than being fixed beforehand.
- Finally, evaluating the architecture on larger and more realistic port scenarios may provide additional insight into whether Global-Local Attention becomes increasingly advantageous when longer-range dependencies are present. Testing on massive terminal topologies where local message passing naturally decays could further solidify the necessity of global communication hubs.

The findings suggest that incorporating global communication mechanisms into reinforcement learning architectures for berth allocation is a promising direction, but that their effectiveness depends strongly on how this information is selected and propagated throughout the scheduling process. Although the observed improvements are accompanied by increased computational requirements, the results demonstrate that representations based on Global-Local Attention provide a statistically significant improvement in performance compared to the GNN approach. This work offers a starting point for future research on attention mechanisms in dynamical scheduling environments and shows their potential to improve the efficiency and reliability of future port operations.

References

[1] T. W. Bank, "Shipping and ports," https://www.worldbank.org/ext/en/topic/transport/shipping-and-ports, accessed: 2026-06-01.

[2] L. Huang, Y. Tan, and X. Guan, "Hub-and-spoke network design for container shipping considering disruption and congestion in the post covid-19 era," *Ocean & coastal management*, vol. 225, p. 106230, 2022, doi: https://doi.org/10.1016/j.ocecoaman.2022.106230.

[3] UNCTAD, "Review of maritime transport 2024," https://unctad.org/publication/review-maritime-transport-2025, accessed: 2026-06-01.

[4] ——, "Maritime vessel turnaround time key performance indicator," https://sft-framework.unctad.org/key-performance-indicator/maritime-vessel-turnaround-time, accessed: 2026-06-20.

[5] Ç. Iris and J. S. L. Lam, "A review of energy efficiency in ports: Operational strategies, technologies and energy management systems," *Renewable and Sustainable Energy Reviews*, vol. 112, pp. 170–182, 2019, doi: https://doi.org/10.1016/j.rser.2019.04.069.

[6] R. Li, X. Zhang, L. Jiang, Z. Yang, and W. Guo, "An adaptive heuristic algorithm based on reinforcement learning for ship scheduling optimization problem," *Ocean & Coastal Management*, vol. 230, p. 106375, 2022, doi: https://doi.org/10.1016/j.ocecoaman.2022.106375.

[7] C. Bierwirth and F. Meisel, "A survey of berth allocation and quay crane scheduling problems in container terminals," *European journal of operational research*, vol. 202, no. 3, pp. 615–627, 2010, doi: https://doi.org/10.1016/j.ejor.2009.05.031.

[8] C. March Moya, F. Schulte, K. Tierney, and N. Yorke-Smith, "A reinforcement learning approach for the dynamic berth allocation problem," not published yet.

[9] I. Beltagy, M. E. Peters, and A. Cohan, "Longformer: The long-document transformer," *arXiv preprint arXiv:2004.05150*, 2020, doi: https://doi.org/10.48550/arXiv.2004.05150.

[10] V.-I. Anica-Popa, "Implementation of global-local attention for encoding," https://github.com/Vlad-Anica/BRP-RL-BAP, 2026, accessed: 2026-06-10.

[11] Y. Lv, M. Zou, J. Li, and J. Liu, "Dynamic berth allocation under uncertainties based on deep reinforcement learning towards resilient ports," *Ocean & coastal management*, vol. 252, p. 107113, 2024, https://doi.org/10.1016/j.ocecoaman.2024.107113.

[12] UNCTAD, "Building capacity to manage risks and enhance resilience," https://resilientmaritimelogistics.unctad.org/, 2026, accessed: 2026-06-20.

[13] A. Lim, "The berth planning problem," *Operations research letters*, vol. 22, no. 2-3, pp. 105–110, 1998, doi: https://doi.org/10.1016/s0167-6377(98)00010-8.

[14] A. Imai, E. Nishimura, and S. Papadimitriou, "The dynamic berth allocation problem for a container port," *Transportation Research Part B: Methodological*, vol. 35, no. 4, pp. 401–417, 2001, doi: https://doi.org/10.1016/s0191-2615(99)00057-0.

[15] S. Kordić, T. Davidović, N. Kovač, and B. Dragović, "Combinatorial approach to exactly solving discrete and hybrid berth allocation problem," *Applied Mathematical Modelling*, vol. 40, no. 21-22, pp. 8952–8973, 2016, doi: https://doi.org/10.1016/j.apm.2016.05.004.

[16] S.-W. Lin, K.-C. Ying, and S.-Y. Wan, "Minimizing the total service time of discrete dynamic berth allocation problem by an iterated greedy heuristic," *The Scientific World Journal*, vol. 2014, no. 1, p. 218925, 2014, doi: https://doi.org/10.1155/2014/218925.

[17] Y. B. Türkoğulları, Z. C. Taşkın, N. Aras, and İ. K. Altınel, "Optimal berth allocation, time-variant quay crane assignment and scheduling with crane setups in container terminals," *European journal of operational research*, vol. 254, no. 3, pp. 985–1001, 2016, doi: https://doi.org/10.1016/j.ejor.2016.04.022.

[18] M. Yu, Y. Lv, Y. Wang, and X. Ji, "Enhanced ant colony algorithm for discrete dynamic berth allocation in a case container terminal," *Journal of Marine Science and Engineering*, vol. 11, no. 10, p. 1931, 2023, doi: https://doi.org/10.3390/jmse11101931.

[19] Z.-H. Hu, "Multi-objective genetic algorithm for berth allocation problem considering daytime preference," *Computers & Industrial Engineering*, vol. 89, pp. 2–14, 2015, doi: https://doi.org/10.1016/j.cie.2015.04.035.

[20] H.-P. Hsu, "A hpso for solving dynamic and discrete berth allocation problem and dynamic quay crane assignment problem simultaneously," *Swarm and Evolutionary Computation*, vol. 27, pp. 156–168, 2016, doi: https://doi.org/10.1016/j.swevo.2015.11.002.

[21] R. M. de Oliveira, G. R. Mauri, and L. A. N. Lorena, "Clustering search for the berth allocation problem," *Expert Systems with Applications*, vol. 39, no. 5, pp. 5499–5505, 2012, doi: https://doi.org/10.1016/j.eswa.2011.11.072.

[22] E. T. Bacalhau, L. Casacio, and A. T. de Azevedo, "New hybrid genetic algorithms to solve dynamic berth allocation problem," *Expert Systems with Applications*, vol. 167, p. 114198, 2021, doi: https://doi.org/10.1016/j.eswa.2020.114198.

[23] S.-W. Lin and C.-J. Ting, "Solving the dynamic berth allocation problem by simulated annealing," *Engineering optimization*, vol. 46, no. 3, pp. 308–327, 2014, doi: https://doi.org/10.1080/0305215x.2013.768241.

[24] E. Lalla-Ruiz, B. Melián-Batista, and J. M. Moreno-Vega, "Artificial intelligence hybrid heuristic based on tabu search for the dynamic berth allocation problem," *Engineering Applications of Artificial Intelligence*, vol. 25, no. 6, pp. 1132–1141, 2012, doi: https://doi.org/10.1016/j.engappai.2012.06.001.

[25] Y. Lv, M. Zou, J. Li, and J. Liu, "Dynamic berth allocation under uncertainties based on deep reinforcement learning towards resilient ports," *Ocean & coastal management*, vol. 252, p. 107113, 2024, doi: https://doi.org/10.1016/j.ocecoaman.2024.107113.