Reinforcement Learning for the Dynamic Berth Allocation Problem

Evaluating the Robustness of a Graph Neural Network Agent under Uncertainty of Estimated Time of Arrival

Lasse Geurts

Supervisors: Carlos March Moya, Neil Yorke-Smith

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
June 19, 2026

Name of the student: Lasse Geurts Final project course: CSE3000 Research Project
Thesis committee: Neil Yorke-Smith, Carlos March Moya, Wendelin Boehmer

Reinforcement Learning for the Dynamic Berth Allocation Problem

Evaluating the Robustness of a Graph Neural Network Agent under Uncertainty of Estimated Time of Arrival

Lasse Geurts
Delft University of Technology
Supervisors: Carlos March Moya, Neil Yorke-Smith

Abstract

The Dynamic Berth Allocation Problem is a port scheduling problem where vessels arrive dynamically over time and must be assigned to a berth. A pre-trained Graph Neural Network (GNN) based reinforcement learning approach solves the problem efficiently but is dependent on estimated times of arrival [11]. However, vessels predominantly arrive later than estimated. this introduces unwanted uncertainty. We evaluate a pre-trained GNN agent under controlled ETA perturbations to create an information gap between the estimated and actual arrival times, using the original full information setting as a reference. The agent is evaluated using a factorial experimental setup over different deviation levels and instances to evaluate robustness under ETA uncertainty. The results show that the agent is robust to optimistic ETA deviations, with only limited performance degradation even at larger introduced deviations. The robustness appears to mainly stem from the reactive nature of the agent and its associated scheduling process, which limits the influence of ETA deviations on the decision making process.

1 Introduction

The schedule of a port depends vessel departure and arrival times, which are uncertain and can be difficult to predict [8]. The efficient allocation of berths is essential for reducing vessel waiting times and operational costs associated with port operations. When vessels are not able to be assigned to a berth immediately upon arrival they might have to wait at anchorage or in virtual queues, which may lead to delays propagating throughout port operations. Since international trade depends on maritime transport for over 80% of global trade volume [13], reducing these delays and improving overall efficiency and robustness of port scheduling is still an immense operational challenge.

A well known optimization problem in port logistics is the Berth Allocation Problem (BAP) [2]. The BAP consists of assigning vessels to berths, whilst taking constraints such as berth availability and berth capacity into account. Additionally, berth-dependent handling times and vessel priorities are integrated into the objective function to evaluate the quality of a schedule. The Dynamic Berth Allocation Problem (DBAP) is an extension of the BAP where vessels arrive over time instead of all being available at the start of the problem [2]. The dynamic nature of DBAP makes the setting more challenging, since decisions are made sequentially whilst still taking into account future arrivals and berth availability.

Recent work by Moya et al. [11] has proposed a Graph Neural Network (GNN) based reinforcement learning agent for solving the DBAP. We evaluate the same pre-trained agent without training or further optimization. This approach allows the agent to learn by interacting with simulated problem instances using a graph-based approach, capturing the relationships between the vessels and the berths. However, the performance of this agent is dependent on the information that is available to it at the time of solving. In particular, the agent relies on estimated times of arrival. These estimated times are reported to a port some time in advance and often differ from the actual arrival times [4]. In practice, the reported times are optimistic as nearly 40% of vessels arrive at least a day later than reported [14]. These inaccuracies may occur due to weather conditions, congestion, operational delays or other disruptions [4].

This potentially creates an information gap between the estimated arrival times known to the agent and the actual arrival times observed in the simulated instances. A policy that performs well under perfect conditions might not necessarily perform well when estimated time of arrival deviations are introduced, especially in highly congested situations where small timing differences can affect future decision making.

In this work, we investigate how the performance of the trained GNN agent for the DBAP is influenced when introducing ETA deviations. To do so, controlled perturbations of the actual arrival times are introduced to create the estimated arrival times. We then evaluate the agent under different congestion levels to assess the robustness under varying operational conditions. We find that the GNN agent is mostly insensitive to small ETA deviations, where larger perturbations introduce small bounded degradations in performance. The effect of the perturbations is non-linear and varies slightly based on the deviation levels and perturbation types used.

The main contributions of this work are as follows:

1. We propose a perturbation framework for introducing optimistic ETA uncertainty into the DBAP, enabling systematic robustness analysis.
2. We evaluate the robustness of a pre-trained GNN agent under different perturbation types, deviation magnitudes and congestion levels.
3. We show that the evaluated GNN agent remains largely robust to optimistic ETA deviations, with only small and bounded performance degredation.

The rest of the paper is structured as follows: Section 2 introduces the background and formal problem description of dynamic the berth allocation problem. Section 3 presents the methodology and experimental setup, with the ETA perturbation framework. Section 4 presents the results of the experiments, which are then discussed in Section 5. Section 6 discusses responsible research considerations and limitations of the paper. Finally, Section 7 concludes the paper and outlines future work.

2 Background and Problem Description

This section introduces the Dynamic Berth Allocation Problem and the graph-based reinforcement learning approach we use. First, the DBAP is formulated as a sequential decision-making problem. Its objective, constraints, reward and state representation are explained. Next, a GNN based approach for solving this problem is introduced. Finally, the information gap that is created between the estimated and actual arrival times when ETA deviations arise is explained.

2.1 Dynamic Berth Allocation Problem

The DBAP is a problem that involves the assignment of arriving vessels to available berths [6]. Unlike the berth allocation problem, where all vessels are present at the start of the evaluation process, vessels in the DBAP arrive dynamically throughout the evaluation process [2]. As a result of this distinction, the assignment process changes from a static to a sequential process where both current and future arrivals are considered [6].

Let $V = \{v_1, v_2, \dots, v_n\}$ be the set of vessels and $B = \{b_1, b_2, \dots, b_m\}$ the set of berths. Each vessel v_i is then associated with an arrival time t_i, a priority value p_i and a list of berth-specific handling times h_{ij} which represents the time needed to service vessel v_i at berth b_j. Finally, let $b(i) \in B$ denote the berth assigned to vessel v_i.

The objective of the DBAP is to minimise operational costs whilst satisfying a set of feasibility constraints. A full mixed integer programming formulation of the DBAP, including sequencing and non-overlap constraints is given in [11]. These constraints ensure that (i) each vessel is assigned to exactly one berth, (ii) no two vessels occupy the same berth at the same time, (iii) the servicing of a vessel cannot start before the arrival time of a vessel.

The objective is to minimise the total priority-weighted cost, which consists of a waiting time and a handling time given as:

$$C_{\text{total}} = C_{\text{waiting}} + C_{\text{handling}} \tag{1}$$

$$C_{\text{waiting}} = \sum_{i=1}^{|V|} p_i \cdot w_i \tag{2}$$

$$C_{\text{handling}} = \sum_{i=1}^{|V|} p_i \cdot h_{i,b(i)} \tag{3}$$

Since there could be significant differences in the number of vessels and berths in the instances, another important metric is the normalized objective cost, defined as:

$$C_{\text{norm}} = \frac{C_{\text{total}}}{|V|} \tag{4}$$

The normalization makes sure that the performance metrics stay meaningful across differently sized instances. Without this normalization, the larger instances would produce larger objective values independent of the scheduling quality of the agent.

The DBAP is formulated as a sequential decision-making problem on a dynamic bipartite graph between vessels and berths [10]. At each decision step t, the state of the environment s_t consists of a set of arrived vessels V_{arr} and a set of future vessels V_{fut} and berth availability B_t.
The action space consists of two actions:

- Assignment actions $a_t = (v, b)$, which allocates vessel $v \in V_{\text{arr}}$ to an available berth in $b \in B_t$. It is important to note that once a vessel has been assigned, it cannot be reassigned.
- Skip action, which advances the simulation in time until either a new vessel arrives or a berth is available again.

The reward is defined as the negative change in the objective value which is normalized by a theoretical lower bound. This converts the cost minimization problem into a reward maximization problem [1]. This ensures that the agent tries to maximize its reward which corresponds to finding the most efficient schedule.

2.2 GNN Reinforcement Learning Approach

We evaluate the robustness of an existing GNN-based reinforcement learning policy proposed for the DBAP [3, 11]. The core architecture of the GNN agent from the original work is left intact. However, the surrounding simulation and arrival-time handling are extended to support experiments regarding ETA uncertainty.

The GNN represents the scheduling environment as a graph with vessel nodes and berth nodes. Nodes encode operational information such as vessel priorities, berth availability, handling times and time-dependent information. Edges represent the relationships between the vessels and berths and capture the assignment dependencies in the scheduling environment.

An important distinction is that future vessels are also included in the graph representation before they physically arrive at the port. As a result the agent can base its decision not only on the currently available vessels but also on estimated future arrivals. This allows the agent to anticipate future congestion and berth utilization, which enables the model to exploit structured representations of future problem elements [10]. However, relying on estimated future arrival times introduces sensitivity to inaccuracies in these estimations.

This approach is different from classical heuristic approaches, like First Come First Served (FCFS), since the heuristic approaches react only to vessels that have already physically arrived [12]. Consequently these heuristics do not use the estimated times of arrival of future vessels, unlike the GNN agent which encodes this information into the state representation.

2.3 Information Gap and ETA Uncertainty

In real life port scheduling, estimated times of arrival may differ from the actual arrival times due to several factors: weather conditions, congestion, operational disruption or routing changes [4]. We consider only optimistic ETA deviations, where actual arrivals are assumed to be later than the estimated arrivals. We study the effect of such ETA inaccuracies on the behaviour of the trained GNN agent.

To evaluate robustness under ETA uncertainty a distinction is made between actual and estimated arrival times [7].

The actual arrival times are used for the physical simulation whereas the estimated arrival times are observed by the GNN for future vessels.

The main experiment introduces optimistic ETA perturbations such that the estimated arrival time is prior to the actual arrival time. Let t_i^{actual} be the physical arrival time of vessel v_i generated by the problem instance. We then model the optimistic ETA uncertainty by sampling a non-negative random variable ϵ_i, where $\epsilon_i \sim \mathcal{E}$ and $\mathcal{E}$ represents an ETA deviation distribution:

$$t_i^{\text{estimated}} = t_i^{\text{actual}} - \epsilon_i, \quad \epsilon_i \geq 0 \tag{5}$$

An alternative formulation that would still ensure optimistic estimated arrival times would be where the actual arrival time is a perturbation of the estimated arrival time:

$$t_i^{\text{actual}} = t_i^{\text{estimated}} + \epsilon_i, \quad \epsilon_i \geq 0 \tag{6}$$

In this case the randomness changes the physical arrival time of the vessels, meaning that t_i^{actual} is no longer fixed in the problem but rather sampled. This would change the underlying congestion pattern of the scheduling problem. By delaying the physical arrival times, vessels may be spread over a wider horizon, reduce queue overlap and lower waiting costs. Therefore changes in performance can no longer solely be attributed to the ETA deviations. For that reason, Equation (5) is used.

Under that formulation vessels are expected to arrive earlier than they physically arrive in the simulation. As a result the future arrivals appear in the graph representation before a berth can actually be assigned. This creates an information gap between the arrivals observed by the simulation and the expectations of the GNN agent, introducing a form of partial observability [7].

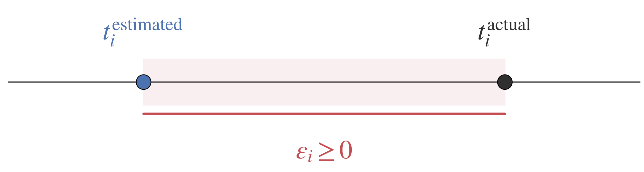

Figure 1: Optimistic ETA formulation showing $t_i^{\text{estimated}} = t_i^{\text{actual}} - \epsilon_i, \quad \epsilon_i \geq 0$, creating the information gap between simulation and reality

3 Methodology and Experimental Setup

This section explains the experimental setup used to evaluate the robustness of the pre-trained GNN agent for the DBAP under ETA uncertainty. The goal of the experiment is not to train but rather to evaluate how robust the pre-trained agent is. First, the trained GNN agent and the heuristic agents are addressed. Second, the changes made to the simulation environment are discussed, which allow the separation of the estimated and actual arrival times such that the agent can be tested under the information gap. Following is the structured process taken to generate the instances on which the experiment is run. Finally, the perturbation framework which defines the type of perturbations used to create the estimated arrival times is introduced.

3.1 Baseline Model

In this paper we evaluate a pre-trained Graph Neural Network reinforcement learning agent as proposed by Moya et al. [11] for the DBAP. The model design is left untouched and no retraining is performed. This agent is evaluated in a fixed setting such that we can isolate the effects of ETA uncertainty on the performance of the model.

To fully evaluate the effect of ETA uncertainty we consider two different evaluation settings applied to the same trained agent. In the full information setting, the agent is evaluated in the originally proposed environment, where the agent is able to see the true arrival times. In the partial information setting, the agent is subjected to the information gap. Importantly, queue updates and vessel activation are always based on the actual arrival times in both regimes. Now the agent no longer sees the actual arrival times, but is only subjected to the estimated arrival times while the actual arrival times stay unchanged. The same underlying instances and conditions are used across both settings to ensure that differences in performances can be based on the structural differences between the information settings.

In addition to the learned GNN agent the classical heuristic agents mentioned in Moya et al. [11] are considered as reference. These heuristics remain unaffected by the ETA perturbations. Unlike the GNN agent, the heuristics do not rely on the estimated times of arrival of the future vessels. They make decisions solely based on the vessels that have already physically arrived at the port, which is exclusively determined by the actual arrival times. As a result, the perturbations applied to create an information gap and the estimated arrival times have no influence on the decision process of the heuristic agents. These agents consequently operate on identical observable states no matter the perturbation type, deviation magnitude or random seed used.

Therefore, heuristic results should strictly be interpreted as a deterministic baseline reference. They provide a stable point of comparison across all perturbation settings, but the heuristics cannot be used themselves as indicators of robustness under uncertainty as their behaviour is fundamentally unaffected by the created information gap.

3.2 Simulation Environment Changes

The simulation environment has been extended by adding an extra attribute to each vessel. Through this addition the environment now maintains both the actual and estimated arrival time of each vessel. This enables the separation of simulation and GNN belief. The simulation logic, including vessel activation, queue updates and scheduling progression is dependent solely on the actual arrival times of all vessels. However, for construction of the state graph the estimated times of arrivals are used, but only for vessels that have not yet arrived. For vessels that have already arrived the nodes are constructed using the actual arrival times, this is due to the fact that arrived ships are known with certainty and represent the current physical state of the port. This separation does not affect the underlying DBAP mechanics, the action space, reward computation or berth assignment rules. The only change that is made is in regard to the information exposed to the agent.

3.3 Instance Generation

Dynamic berth allocation instances were generated artificially using a controlled stochastic process implemented by Moya et al. [11]. The test set configuration was used to guarantee a controlled and reproducible set of problems. Each problem instance consists of a set of vessels and berths, where vessels' inter-arrivals follow an exponential distribution and are controlled via a congestion parameter C. This parameter controls the density of the actual vessel arrival times, where a lower value represents a higher congested port. Handling times and priorities are sampled from uniform distributions as per Moya et al. [11].

The test set used for this experiment uses the parameters laid out in Table 1. For all combinations of these parameters, 216 instances (see Table 2) are generated using fixed random seeds, resulting in stochastically different but reproducible variants of the same instance. Each seed produces a complete problem instance with independently sampled handling times, priorities and arrival patterns. By evaluating the agent over both structural variation via different problem sizes and stochastic variation via different random patterns we get reproducible yet meaningful performance metrics.

Table 1: Instance size and structural parameters.

Number of vessels	$\{80, 100, 120\}$
Number of berths	$\{5, 10, 15\}$
Congestion level C	$\{0.5, 1, 2\}$

Table 2: Stochastic generation settings and total instance count.

Seeds per configuration	$\{42, \ldots, 49\}$ (8 seeds)
Total configurations	$3 \cdot 3 \cdot 3 = 27$
Total instances	$27 \cdot 8 = 216$

3.4 Estimated Time of Arrival Perturbation Framework

To be able to evaluate robustness under ETA deviations we must add controlled perturbations at instance level to the estimated arrival times and thus keeping the arrival times fixed for creating the information gap. The main objective of this perturbation framework is adding systematic and stochastic perturbations in the arrival time information while still keeping the underlying generated problems the same.

As can be seen in Table 1, the generated instances significantly differ in size. To ensure that the perturbations are meaningful compared to the size of the instance the magnitude of the perturbations is normalised to the instance level mean handling time $\bar{h}$ (See Equation (7)). For N vessels and M berths, h_{ij} represents the handling time of vessel i at berth j.

We express $\bar{h}$ as follows:

$$\bar{h} = \frac{1}{N \cdot M} \sum_{i=1}^{N} \sum_{j=1}^{M} h_{ij} \tag{7}$$

In the experiment we consider four types of ETA perturbations, where all magnitudes are expressed as multiples of $\bar{h}$. Each of the perturbation types represents a different structural form of ETA uncertainty. The perturbations are denoted as the ETA-noise ϵ_i defined in Equation (5).

The first perturbation type is a systematic shift, this represents the most simple possible perturbation model where all vessels have a constant delay β that is applied to all vessels. This is defined as:

$$\epsilon_i = \beta, \quad \beta \geq 0$$

This perturbation type models systematic bias in the estimated times of arrival or port-wide disruptions that shift the full schedule of the port. While this type of perturbation is structurally simple, it isolates the global temporal misalignment between perceived and actual system dynamics.

The second type of perturbation considered is positive-mean Gaussian noise to model vessel specific uncertainty, a common assumption in stochastic and forecasting error models. [5] [9]:

$$\epsilon_i \sim \mathcal{N}(\mu, \sigma^2),$$

which is then truncated to ensure optimistic errors, so $\epsilon_i \leftarrow \max(0, \epsilon_i)$. This perturbation type captures the stochastic deviations from the optimistic ETA reports. The positive mean makes sure that the model aligns with the assumption that the estimated times of arrival are generally optimistic.

To reflect asymmetric delay distributions commonly observed in maritime logistics we introduce exponential delay, a standard modelling choice for non-negative assymetric stochastic processes [9]:

$$\epsilon_i \sim \text{Exponential}(\lambda)$$

The non-negative nature of the Exponential distribution conforms with our underlying assumption of optimistic errors. Since the distribution introduces right-skewed variability this allows us to model the scenario where most of the vessels have a small delay whereas a smaller subset is heavily delayed.

Finally, we consider correlated uncertainty which is a combined structure of global and vessel-specific uncertainty:

$$\epsilon_i = \eta_{global} + \eta_i$$

where η_{global} represents port-wide disruption and η_i represents the uncertainty for specific vessel v_i. This type is meant to model realistic port conditions where both local and global disruptions intertwine to create larger delays.

As previously mentioned, the magnitude of uncertainty is relative to $\bar{h}$. We define a set of scalars $\mathcal{D}$, expressed as multiples of $\bar{h}$, to represent the deviation levels:

$$\mathcal{D} = \{\alpha \cdot \bar{h} \mid \alpha \in \{0, 0.0625, 0.125, 0.25, 0.5, 0.75, 1, 1.5, 2\}\}$$

Here, $\mathcal{D}$ defines the absolute scale of perturbation applied to any vessel in a specific instance. The case when $\alpha = 0$ represents the case where there is no perturbation, so where estimated and actual arrival times are the same. Thus, in this setting both the information regimes receive identical state information. This serves as a consistency check between the full and partial information agents, ensuring that the perturbation framework reproduces the original environment and that the evaluation across both information regimes is being executed on the same underlying instances.

3.5 Experiment Procedure

The evaluation of the experiment is performed in a controlled factorial design where both information regimes are evaluated over identical problem instances with systematically varying ETA perturbations. By ensuring both regimes see the same problem instance we can attribute the performance differences solely to the asymmetry of information instead of structural differences in the problem instance.

For each generated instance, the mean handling time $\bar{h}$ is computed to determine the scale of the instance. Subsequently a perturbed variant of the originally generated instance is generated based on the selected perturbation type, deviation level $d \in \mathcal{D}$, and a random seed $s \in \{42, \ldots, 47\}$. This perturbed instance is then independently evaluated across all different agents where each information regime receives an identical copy of the perturbed instance. For the stochastic perturbation types (Gaussian, Exponential and correlated noise) multiple seeds are used to take random variability in the sampling process into account. For the deterministic perturbation types, namely systematic shift and the case $\alpha = 0$ we evaluate using a single seed, since no randomness needs to be accounted for.

After each evaluation run we store performance metrics such as waiting time, handling cost and total cost are computed alongside some behavioural metrics deduced from the action logs created by the agents.

3.6 Statistical Reporting

The primary goal of the evaluation is to evaluate how the GNN-agent performs under ETA uncertainty. To systematically study this effect, the deviation magnitude set $\mathcal{D}$ introduced in Section 3.4 is used to control the strength of perturbations applied to the actual arrival times. Each element $\alpha \in \mathcal{D}$ is a distinct perturbation strength expressed as a multiple of $\bar{h}$

Results are therefore displayed as functions of α, which allows direct comparison of robustness and performance across different uncertainty levels. This gives a consistent scale across instances of different sizes, making sure that all effects can be attributed to the perturbations.

All the data points across the $\mathcal{D}$-axis are computed over all problem instances, perturbation types and random seeds. This gives results based on repeated stochastic evaluations per configuration, allowing results to be interpreted as population-level data instead of single outcomes. By doing this we can systematically compare the results between the perturbation types whilst keeping the interpretation centred on the performance over the $\mathcal{D}$-axis

4 Results

This section discusses the results of the performed ETA perturbation experiments over different levels of α and over different structural types of deviations. First, the section illustrates the overall robustness of the agent to uncertainty of estimated arrival times. Second, the impact of the different structural types of deviation is shown. Finally, instance characteristics such as operational load and structural capacity are studied for their effects on the robustness of the agent.

All results, except those in Figure 4, are reported as changes relative to a baseline set by the full information regime. This ensures that all the changes can be attributed to the ETA information gap observed by the partial information regime, rather than structural differences in the instances or model configurations. Therefore the results focus on how increasing levels of uncertainty α affect the robustness of the partial information regime.

4.1 Overall Robustness to Estimated Time of Arrival Uncertainty

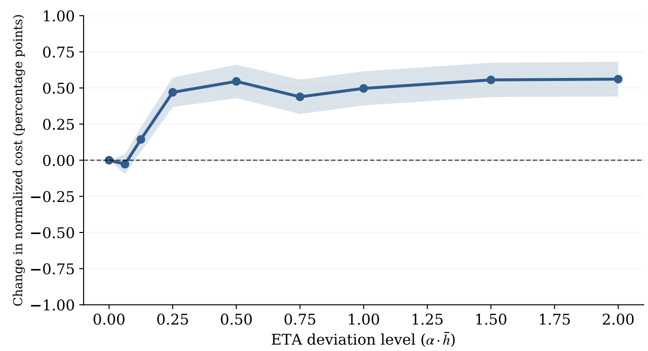

Figure 2: The mean percentage point in the normalised cost compared to the full information regime as a function of the deviation level α. The shaded region indicates the 95% confidence interval of the mean across the instances.

Figure 2 displays the mean percentage point change in performance for all the perturbation types as a function of the deviation level α. Across all the perturbation types, the performance of the GNN agent displays consistent non-linear behaviour as an effect of increasing the ETA uncertainty. At smaller α ($\alpha \leq 0.125$), the performance is relatively indistinguishable from the baseline with only minor fluctuations observable. This indicates near-complete robustness in regard to small optimistic errors in the estimated times of arrival.

A sharp increase of the normalised cost is between $\alpha = 0.125$ and $\alpha = 0.25$ where the mean percentage point nearly triples, this suggests that there is some kind of transition point in which the model becomes increasingly sensitive to ETA perturbations. Beyond $\alpha = 0.25$, the robustness solidifies with the mean percentage change stabilising as it fluctuates in a narrow band. Important to note is that after $\alpha = 0.5$ the function is no longer monotonic as the increasing deviation levels no longer lead to an increased normalised cost leading to a plateau.

The standard error of the mean (SEM) does increase gradually with the magnitude of the added uncertainty, indicating higher variability for the instances that were strongly perturbed. However, this increase is very small and seems to be bounded, suggesting that the perturbations do not dominate the decision making of the agent.

The performance of the heuristics does not change by design, as mentioned in Section 3.1. Their performance is merely a fixed reference point as all observed performance changes in Figure 2 solely reflect the changes in the GNN agent. Consequently, any changes in the relative performance gap between the GNN and the heuristics can be directly seen as the loss of performance of the GNN as displayed in Figure 2

4.2 Effect of Perturbation Type on Robustness

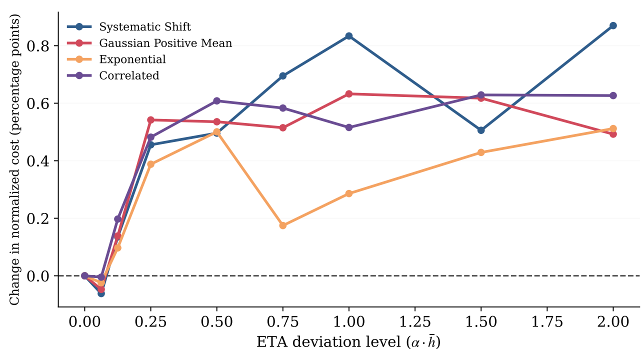

Figure 3: The mean percentage point in the normalised cost compared to the full information regime as a function of the deviation level α. Each curve corresponds to a different perturbation type indicated by the legend

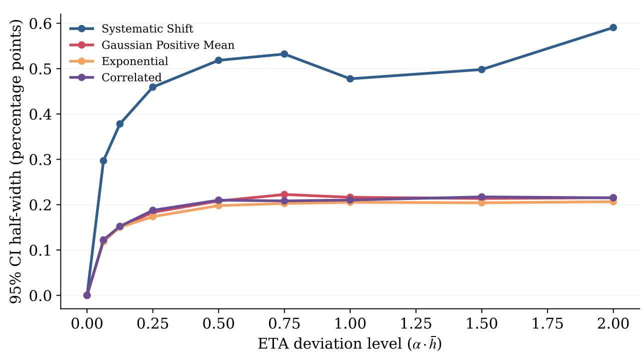

Figure 4: The uncertainty in performance across the perturbation types displayed as 95% confidence intervals.

Figure 2 captures the overall sensitivity to ETA perturbations; however, it does not show the differences between the types of perturbations introduced in Section 3.4. Thus, we analyse how the different structural forms of perturbations influence the performance robustness under increasing levels of deviation α. Figure 3 shows the mean percentage point change per perturbation type across the deviation levels. Figure 4 complements this by displaying the associated uncertainty of each perturbation type, reported as 95% confidence intervals. The different perturbation types follow the general trend established in section 4.1, however the profiles of the response stability are different.

Systematic shift shows the most uncertain response pattern. The performance changes are sensitive to the deviation magnitude α, but do not follow a smooth or consistently monotonic pattern. Instead the response is characterised by higher fluctuations over the deviation levels α as can be seen in Figure 4. This suggests that global temporal misalignment introduces effects that cascade throughout scheduling process, where early errors in decision making propagate through the system in a non-linear way.

In contrast, the stochastic processes such as the positive mean Gaussian and the exponential distributions produce more regular response curves. The Gaussian despite minor fluctuations has a mostly consistent increase in performance degradation over greater magnitude deviations. The exponential distribution follows much of the same pattern, although it does generally have a more contained response to the perturbations. This suggests that asymmetric but independent perturbations have a more predictable overall effect. For both distributions the uncertainty remains at a relatively lower level, indicating that vessel specific noise is partially absorbed by the scheduling policy.

The correlated perturbation type has the most stable performance in terms of the mean percentage point changes in Figure 3. Compared with the other perturbation types it has smaller fluctuations indicating that the agent acts more predictably when the perturbations contain both global and vessel specific components. While the performance curve is more stable, we do not see this translate into a lower uncertainty in the results of Figure 4. The confidence intervals associated with the correlated perturbations are very similar to the ones introduced by the stochastic models. Consequently, correlated uncertainty does not necessarily make the outcome less variable but instead has a more consistent deterioration in performance as deviation level α grows.

4.3 Impact of Instance Characteristics on Robustness

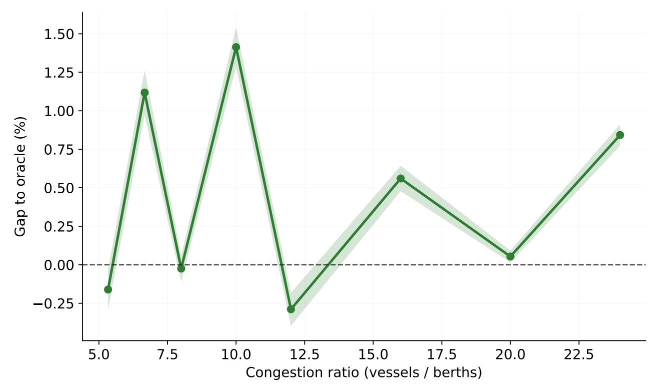

Figure 5: The impact of the amount of vessels per berth on the relative change in normalised cost compared to the full information regime. Each point represents the mean performance across instances that share the same amount of vessels per berth.

To analyse whether the structural properties of the generated instances influence the robustness of the agent, two

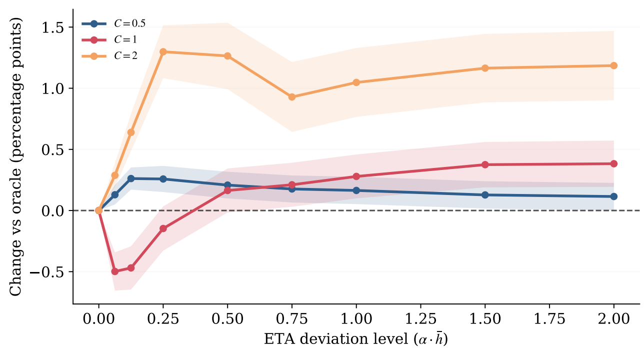

Figure 6: Sensitivity of the relative change in normalised cost compared to the full information regime to ETA deviations under different levels of congestion. Each curve represents a congestion level on which an instance was generated and points represent the mean outcomes across these instances.

instance-level variables are analysed: the berth-vessel ratio and the congestion parameter $c \in C$. These two variables capture two aspects of instance difficulty, static capacity and operational load.

Figure 5 displays the relationship between the ratio of vessels per berth and the mean percentage point change in the normalised cost. No monotonic relationship can be found between the ratio and the degradation of performance, instead the effect of the berth vessel ratio is highly non-linear as across the deviation levels alternating positive and negative deviations can be found.

In particular, the lower ratios do not outperform the higher ratios, and no consistent ordering can be established over the tested values. This suggests that the structural capacity of a port is not a strong predictor of the robustness of the GNN agent. A good thing to notice is that although differences exist across the ratios the magnitude of the relative change of mean normalised cost is relatively low, indicating that although differences exist they do not dominate the performance.

Figure 6 displays the impact of the congestion parameter as defined in Section 3.3. In contrast to the berth-vessel ratio a stronger pattern emerges in this parameter. For the two more congested levels ($C = \{0.5, 1\}$) performance differences between the levels are minimal and do not have a consistent natural ordering, with congestion level 0.5 not outperforming congestion level 1 consistently by having a lower increase in the change in normalised cost. However, when looking at congestion level 2, we can see there exists a separation between the lower and higher levels. In particular, the lowest congestion level consistently leads to significantly larger performance deviations.

Since the relationship is not strictly monotonic over the congestion levels, the congestion levels seem to act more like a threshold factor. Once significant enough levels of congestion are reached it no longer becomes a pronounced factor in the robustness of the agent. This indicates that the congestion parameter does not independently determine the performance but rather modulates the impact of ETA deviations.

Overall, the results indicate that the GNN agent exhibits generally stable but non-uniform sensitivity to ETA uncertainty. Performance stays close to the baseline for smaller perturbations, for larger perturbations the performance has small but bounded degradation which stabilises at higher uncertainty levels. The differences between the perturbation types indicates that the type of noise introduced has an influence on the form of degradation, with systematic shift and stochastic errors affecting the performance in separate but comparative ways. Additionally, instance-level properties like the congestion parameter have a more pronounced effect, indicating that operational load plays a stronger role in affecting the performance under ETA uncertainty than structural capacity does.

5 Discussion

The results indicate that ETA perturbations have little influence on the performance of the GNN agent, even under relatively large deviations. Whilst Section 4 assesses this effect empirically, this section interprets the observed robustness in terms of the underlying structure of the GNN agent and simulation environment. In particular, we analyse why uncertainty in estimated arrival times has little to no impact on the decision making and what this tells us about the use of anticipatory reasoning by the agent.

An important factor in the robustness we observe is the fact that the agent has no way to make a commitment to a future vessel by reserving a berth for a vessel. The assignment decisions are made solely based on the set of vessels that have already physically arrived and assigning a vessel to a berth does not require or assume future capacity reservations. This means that the information of future vessels does not put a constraint on the current decision making in a strict sense, it only provides contextual information about what may arrive later on.

Because of this the errors introduced in the estimated arrival times do not provide a meaningful enough distortion on the agent's decision process. Hence, even when there is a shift in the estimated arrival times, the relative attractiveness of all the feasible actions the agent can execute is mostly unchanged, since the ETA perturbations do not actually alter the set of feasible actions. In other words, the policy executes a regime which is mostly local in time and does not depend on long term allocation of the future vessels.

The lack of such a long term allocation system plays a crucial part in the robustness to ETA uncertainty. In systems where making a commitment to future vessels is possible, the ETA errors can propagate throughout the system and make the system sensitive to uncertainty. However, the evaluated GNN agent only assigns vessels to berths when those vessels are physically present, preventing the forecasting errors from locking into suboptimal future decisions.

As a result, the perturbations have a very limited way to exert influence on the performance of the agent. The agent is presented with distorted information about the estimated times of arrival, but this does not translate to irreversible scheduling decisions or structural constraints placed on future decisions. This explains why even the relatively large deviations in the estimated arrival times lead to only minor changes in performance.

Overall, the results of the experiment suggest that the robustness to ETA deviations is not because of sophisticated anticipation of the future arrivals, but rather an effect of the inherently reactive nature of the underlying scheduling process. However, the variation observed between the different congestion levels indicates that the robustness is not uniform and can be modulated by instance-level characteristics.

6 Responsible Research

The experiment of this paper is done entirely within a simulated environment of the DBAP and does not involve any human participants, personal data or operational port systems. Therefore, no privacy-sensitive or ethically restricted data is used. The objective is purely analytically based, focussing on the performance of algorithms under perturbations instead of real-world decision processes.

From an experimental design standpoint, fair and unbiased comparisons are ensured throughout the evaluation process. All agents are exposed to the identical perturbed instances ensuring that all performance deviations can be attributed to the perturbations. Furthermore, the GNN-agent is used in inference mode, ensuring no form of implicit tuning happened during evaluation.

The evaluation framework and results are reproducible. All the problem instances are generated using reported strategies and reported fixed random seeds, the perturbation levels are specified through set $\mathcal{D}$ and the predefined probability distributions. Stochastic effects are accounted for through repeated evaluation using different random seeds, making variability explicit rather than hidden.

While the experiment is designed to reflect real world uncertainty in estimated arrival times, the perturbation types used remain a simplification of the real-world conditions. Therefore the results should be interpreted as evidence of robustness under controlled uncertainty rather than exact operational performance in real life ports

7 Conclusions and Future Work

In this paper, we evaluated the robustness of a pre-trained Graph Neural Network reinforcement learning agent designed for the Dynamic Berth Allocation Problem under uncertainty of the Estimated Times of Arrival. To find the impact of ETA uncertainty, an information gap was introduced between the estimated times of arrival known to the agent and the actual arrival times used for the physical simulation of the instance. By keeping the underlying instances identical and perturbing only the information known to the agent, we can attribute the effects on the performance of the agent solely to the introduced perturbations.

The results show that the agent is robust to the introduced optimistic ETA deviations. Across all the perturbation types and deviation magnitudes, only a very limited performance degradation is observed. Additionally, the structure of the introduced perturbations proved to have limited influence. A systematic shift of the estimated times of arrival showed to have the most variable response patterns, whereas the stochastic and correlated perturbations had more consistent results. The instance level characteristics showed to have minor influence as well, with congestion having a stronger influence on the degradation than the vessel to berth ratio.

These findings suggest that the robustness is not primarily the result of sophisticated reasoning of future arrivals, but rather due to the structural design of the decision making process of the agent. Vessels can only be assigned to a berth once they physically arrive at the port and the agent is unable to make a commitment to a future arriving vessel by reserving a berth for it. As a result, the ETA information from future arriving vessels is mainly used as contextual information rather than a basis for long-term planning. Consequently, the introduced ETA deviations can only work a limited influence on the set of feasible actions and do not significantly alter the relative attractiveness of the feasible actions. This anticipatory reasoning approach thus naturally limits the impact of ETA uncertainty on the performance of the agent.

Future work could investigate a scheduling policy that does include mechanisms that allow for explicit berth reservations for future vessels or different future commitment mechanisms. These mechanisms allow the agent to widen its planning horizon by reserving berths for future vessels. If ETA predictions are sufficiently accurate, this additional anticipatory capability could improve berth utilisation, reduce waiting times and lead to overall lower cost associated with the created schedules. However, these same mechanisms are likely to increase the sensitivity to ETA deviations, as inaccurate predictions now lead to inaccurate reservations. Studying the trade-off between the use of anticipatory planning and overall robustness to ETA uncertainty would provide insight into what extent future arrival information can be used in berth allocation problems.

Generative AI Disclosure

ChatGPT 5.3 was used exclusively for language editing, clarification of technical writing and LaTeX-formatting in this paper. An AI-based coding assistent with autocomplete from IntilliJ IDEA was used as a programming help for code completion, writing comments/documentation and developmental support. All research design, implementation, analysis and validation were carried out by the author, who assumes full responsibility for the work.

References

[1] Irwan Bello, Hieu Pham, Quoc V. Le, Mohammad Norouzi, and Samy Bengio. Neural Combinatorial Optimization with Reinforcement Learning. 2017. doi: 10.48550/arXiv.1611.09940.

[2] Christian Bierwirth and Frank Meisel. A survey of berth allocation and quay crane scheduling problems in container terminals. *European Journal of Operational Research*, 202(3):615–627, 2010. ISSN 03772217. doi: 10.1016/j.ejor.2009.05.031.

[3] Quentin Cappart, Didier Chételat, Elias B. Khalil, Andrea Lodi, Christopher Morris, and Petar Veličković. Combinatorial Optimization and Reasoning with Graph

Neural Networks. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, pages 4348–4355, 2021. ISBN 978-0-9992411-9-6. doi: 10.24963/ijcai.2021/595.

[4] Zhong Chu, Ran Yan, and Shuaian Wang. Vessel arrival time to port prediction via a stacked ensemble approach: Fusing port call records and AIS data. *Transportation Research Part C: Emerging Technologies*, 176: 105128, 2025. ISSN 0968090X. doi: 10.1016/j.trc.2025.105128.

[5] Rob J Hyndman and George Athanasopoulos. *Forecasting: principles and practice*. OTexts, 2018. ISBN 0987507117.

[6] Akio Imai, Etsuko Nishimura, and Stratos Papadimitriou. The dynamic berth allocation problem for a container port. *Transportation Research Part B: Methodological*, 35(4):401–417, 2001. ISSN 01912615. doi: 10.1016/S0191-2615(99)00057-0.

[7] Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1-2):99–134, 1998. ISSN 00043702. doi: 10.1016/S0004-3702(98)00023-X.

[8] Jana Ksciuk, Stefan Kuhlemann, Kevin Tierney, and Achim Koberstein. Uncertainty in maritime ship routing and scheduling: A Literature review. *European Journal of Operational Research*, 308(2):499–524, 2023. ISSN 03772217. doi: 10.1016/j.ejor.2022.08.006.

[9] Averill M.. Law and W. David. Kelton. *Simulation modeling and analysis*. McGraw-Hill, 1991. ISBN 0070366985.

[10] Nina Mazyavkina, Sergey Sviridov, Sergei Ivanov, and Evgeny Burnaev. Reinforcement learning for combinatorial optimization: A survey. *Computers & Operations Research*, 134:105400, 2021. ISSN 03050548. doi: 10.1016/j.cor.2021.105400.

[11] Carlos March Moya, Frederik Schulte, Kevin Tierney, and Neil Yorke-Smith. A Reinforcement Learning Approach for the Dynamic Berth Allocation Problem. In *LOGMS26*, 2026.

[12] Michael. Pinedo. *Scheduling: theory, algorithms, and systems*. Prentice Hall, 2002. ISBN 9780130281388.

[13] United Nations Trade and Development. *Review of Maritime Transport 2025*. United Nations, 2025. ISBN 9789211591774. doi: 10.18356/9789211591774.

[14] Bert Vernimmen, Wout Dullaert, and Steve Engelen. Schedule Unreliability in Liner Shipping: Origins and Consequences for the Hinterland Supply Chain. *Maritime Economics & Logistics*, 9(3):193–213, 2007. ISSN 1479-2931. doi: 10.1057/palgrave.mel.9100182.