e-Navi

Available online at http://www.e-navigation.kr/
e-Navigation Journal

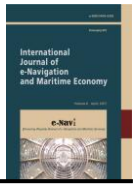

Original article

Improving the Realism and Carbon-Reduction Performance of Weather Routing Using AIS-Derived Maritime Traffic Networks

Ung-Gyu KIM[a], B. Gong HWANG[b], G. Hyun KIM[c], J. Soo KIM[d]

[a] R&D Center, MAPSEA Corp, Korea, ugkim@mapseacorp.com

[b] R&D Center, MAPSEA Corp, Korea, bk22106@mapseacorp.com, Corresponding author

[c] R&D Center, MAPSEA Corp, Korea, kimgh@mapseacorp.com

[d] R&D Center, MAPSEA Corp, Korea, jkim@mapseacorp.com

Abstract

Efficient yet realistic ship routing is critical for reducing fuel consumption and greenhouse-gas emissions. However, conventional weather-routing algorithms often produce mathematically optimal routes that conflict with the paths mariners use. This study presents a hybrid approach that constrains physics-based weather routing within an AIS-derived maritime traffic network (MTN) built from one year of global Automatic Identification System data. The MTN represents common sea lanes as a graph of approximately 10,956 waypoints (nodes) and 17,561 directed edges. Using this network, an optimal low-emission route is computed via graph search and then compared against both a traditional unconstrained route and an advanced weather-routing model (VISIR-2). In a May transition-season case (Busan–Singapore voyage), the AIS-constrained route reduced fuel consumption and CO_2 emissions by about 1.9% relative to the fastest feasible route, while closely following real traffic corridors (over 90% overlap with actual 2024 AIS tracks). While this 1.9% saving does not reach the high-end potential of an unconstrained, state-of-the-art model like VISIR-2 (which can demonstrate double-digit savings in certain conditions), it is achieved with an increase in transit time of ~6.5 h (≈3.2%). This represents a crucial trade-off, prioritizing operational realism and adherence to real-world traffic corridors over maximum theoretical efficiency.

Keywords: AIS-derived Maritime Traffic Network (MTN); Weather-Optimized Ship Routing; Fuel and CO_2 Emission Reduction; Graph-Based Path Optimization; Trajectory Clustering; Douglas–Peucker Algorithm

https://doi.org/10.52820/j.enavi.2025.24.110

1. Introduction

Maritime transport handles about 90% of global trade by volume and contributes roughly 3% of global CO_2 emissions, making efficient voyage planning crucial for both economic and environmental reasons. Traditionally, ship routes are planned manually by experienced navigation officers using nautical charts, sailing directions, and weather forecasts. This process is time-consuming and heavily reliant on the mariner's local knowledge. In unfamiliar waters, even seasoned officers may lack insight into preferred lanes and best practices, which can lead to suboptimal or unsafe routes.

Modern weather-routing algorithms can automatically optimize a ship's path for objectives like minimum transit time or fuel consumption under forecast ocean conditions. However, these physics-based optimizers (e.g., isochrone or grid-based models) often neglect the collective wisdom embedded in historical routes. As a result, purely optimized paths may cut through areas that navigators typically avoid (for example, entering busy separation schemes or shallow waters), or take unconventional tracks that—while mathematically optimal—conflict with real-world operational constraints. Mariners must then adjust or ignore such routing advice, undermining its utility. There is a clear need for routing solutions that balance efficiency with realism.

Recent developments make this integration both feasible and timely. Abundant AIS data now allow data-driven extraction of maritime route networks, while the drive for decarbonization in shipping (e.g., the IMO's GreenVoyage2050 initiative suggests weather routing can save around 5% fuel) creates motivation to improve voyage efficiency. In essence, the opportunity exists to merge two research threads: (1) AIS data mining to capture typical route networks, and (2) physics-based optimization for weather routing. By merging these approaches, the collective intelligence of historical ship navigation can be exploited to guide the optimizer—yielding route plans that a ship's master would find familiar and credible, while still achieving significant reductions in fuel consumption and CO_2 emissions.

Contributions: This paper introduces a novel data-informed weather routing framework that addresses the above challenge. The key contributions are as follows:

Hybrid routing methodology: A routing algorithm is formulated that integrates an AIS-derived MTN with classical weather routing. By constraining the search space to realistic, well-travelled pathways, the method produces routes aligning with actual maritime traffic patterns while still optimizing fuel efficiency under given weather conditions.

Case study demonstration: In a major shipping lane (Busan–Singapore), a representative route network is constructed from historical AIS data, and the hybrid routing approach is demonstrated. The optimized route on this network reduces fuel consumption and CO_2 emissions by approximately 1.9% compared to a conventional unconstrained route, in exchange for a 3.2% increase in travel time. While the AIS-constrained route does not achieve the same level of emissions reduction as VISIR-2's fully optimized fuel-minimizing path, it provides significant savings while adhering closely to established traffic corridors (over 90% overlap with real AIS tracks), enhancing operational realism.

Benchmarking with state-of-the-art: The proposed method is benchmarked against the state-of-the-art VISIR-2 model and a standard grid-based algorithm. This comparison highlights the trade-off between our realism-focused approach and purely mathematical optimization. While VISIR-2 can achieve significantly higher fuel savings (up to 49% in some scenarios) by taking unconstrained, and often circuitous, routes, our method delivers consistent savings while ensuring high compliance with established sea lanes and operational practices.

2. Background and Related Work

AIS-Derived Maritime Traffic Networks: Rather than treating the ocean as a blank canvas, an AIS-derived MTN provides a discrete set of waypoints and "roads" inferred from where ships travel. In the past decade, researchers have explored various techniques to construct such networks from AIS trajectory data. Early works focused on visualizing the density of ship traffic. For example, Lee and Cho (2022) applied kernel density estimation to identify major shipping lanes in Korean waters, distinguishing main routes from branch spurs. Density maps can reveal traffic hotspots but do not directly produce a connected graph of routes for path planning. Later, algorithmic approaches emerged: Zhang et al. (2018) used a line simplification algorithm

(Douglas–Peucker) to reduce AIS trajectories and then identified representative routes, addressing noise and minor detours in raw tracks. More recently, Onyango et al. (2022) formalized a three-step pipeline for AIS-based network extraction: detect significant course-change points, cluster them into waypoints, and connect waypoints into a graph. This yields a simplified maritime route network reflecting typical navigation patterns. Such studies demonstrate that AIS data can indeed be distilled into a graph structure that preserves essential routes. However, most were motivated by navigation safety or anomaly detection—e.g. identifying outlier voyages or common patterns for decision support—rather than directly integrating with weather routing for fuel optimization.

Weather Routing and Emissions: Marine weather routing is the process of finding an optimal path for a ship given forecast of winds, waves, currents, and sometimes other factors (like ice). Traditional algorithms (e.g. the isochrone method or grid-based dynamic programming) aim to minimize travel time and can be extended to minimize fuel by accounting for added resistance from waves or reduced speed in heavy weather. VISIR-2 is an example of a state-of-the-art model that uses graph search on a high-resolution grid to compute least-time or least-risk routes accounting for wave and current effects. Recent research has also looked at routing for emissions: for instance, Mannarini and Carelli (2019) estimated the carbon-intensity savings of optimized routes, and a new VISIR-2 module (Mannarini et al., 2024) allows direct computation of least-CO_2 paths. Other work has incorporated specific operational constraints into routing—e.g. adjusting routes to avoid dangerous resonant roll conditions in rough seas by Mannarini et al., (2016)—but these approaches still optimize in a continuous space without using historical route data. The International Maritime Organization has identified weather routing as a key measure for emissions reduction, projecting modest fuel savings (on the order of a few percent) for typical voyages. To date, however, the integration of historical traffic patterns into automated weather routing remains largely unexplored in the literature.

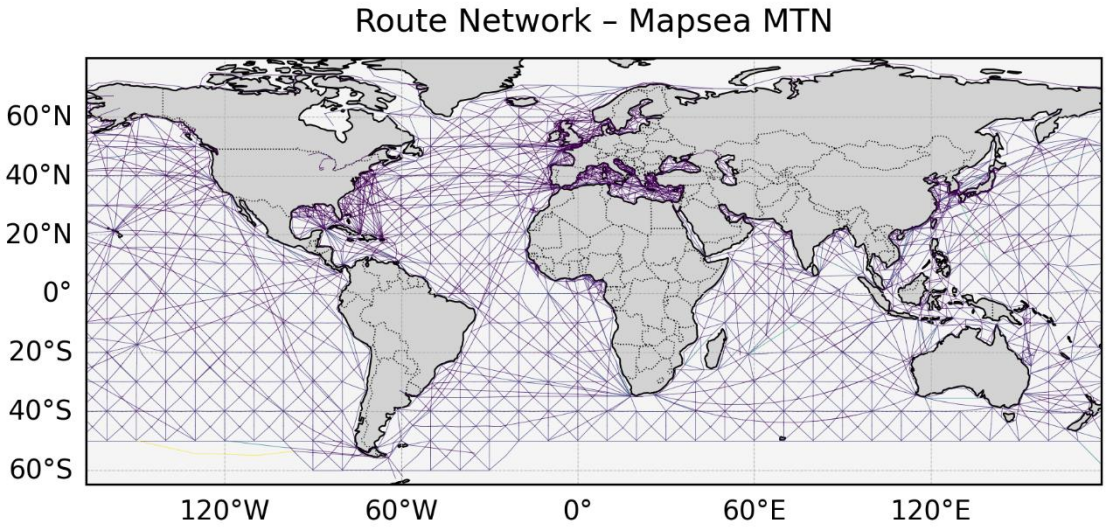

Figure 1: Global AIS-Derived Maritime Traffic Network (Mapsea MTN). The network contains 10,956 nodes and 17 561 edges, colour-coded here by edge density

3. Methodology and Case Study Setup

Following the quasi-intelligent maritime-route-extraction framework proposed by Onyango et al. (2022), the MTN used for routing is extracted from historical AIS data via a three-step process: (i) trajectory simplification and manoeuvre-point detection, (ii) waypoint discovery using Hierarchical Density-Based Spatial Clustering of Applications with Noise (HDBSCAN), and (iii) edge construction through graph theory. Except for the hyper-parameter values discussed later, the implementation strictly adheres to the algorithms and workflow described in that study. One year of global AIS records (January–December 2024) was processed to derive the network, focusing on voyages between the chosen origin and destination. The detailed procedure (summarised in

Figure 1) therefore mirrors Onyango et al.'s original approach and is only adapted here to a global scale and to the Busan–Singapore corridor considered in our case study.

3.1 AIS-Derived Graph Construction:

One year of global AIS records (January–December 2024) was processed to derive the network, focusing on voyages between the chosen origin and destination. The steps include trajectory simplification, maneuver point detection, waypoint clustering, and graph generation, as detailed below.

Trajectory Simplification: Raw AIS trajectories can contain thousands of points with minor zigzags or drift. The Douglas–Peucker (DP) line simplification algorithm was applied to each voyage track to remove extraneous points while preserving the overall route shape. By using a tolerance on the order of a few kilometers, the key turning points of each path (e.g. major course changes around headlands or straits) are retained while small deviations are discarded. The result is a simplified polyline for each voyage, capturing its characteristic waypoints (origin, destination, and significant bends). This denoising focuses attention on meaningful maneuvers. (The DP tolerance value was chosen to balance geometric fidelity against graph sparsity.)

Maneuver Point Detection: The simplified polylines were next analyzed to identify significant maneuver points—typically the vertices of the simplified tracks where ships changed course notably (for example, entering the Taiwan Strait or rounding a cape). These points represent candidate waypoints that many voyages have in common.

Clustering Waypoints (HDBSCAN): The collected maneuver points from all voyages were then clustered to find geographic locations that serve as common waypoints. A density-based clustering algorithm, HDBSCAN, was used to allow discovery of route clusters of varying traffic density. The HDBSCAN hyperparameters were set to min_cluster_size = 50 and min_samples = 10, meaning each cluster (waypoint) had at least 50 contributing voyages. This yielded a set of representative waypoints—each corresponding to a navigationally significant location where ships tend to converge or turn. A minimum cluster size of 50 was chosen to filter out spurious one-off turns while retaining important waypoints. In fact, a sensitivity test varying to this threshold (±20) showed negligible impact on the overall network structure or routing results, indicating the clustering is robust to this parameter.

Graph Construction: Finally, a directed graph was constructed where the nodes are the identified waypoints and edges represent frequently traveled route segments between them. Two waypoints were connected by an edge if a substantial number of AIS voyages followed that segment in sequence. Direct connections were also added for obvious port-to-port legs (e.g. from Busan directly to a common junction waypoint) to ensure network connectivity, even if few voyages took a perfectly straight line. Each edge was assigned a geodesic distance (great-circle distance between the waypoints) and later an environmental cost (travel time or fuel for given conditions). The resulting global MTN contains $|V|$ = 10,956 nodes and $|E|$ = 17,561 directed edges (spread across 41 connected components; mean node degree = 3.21, max degree = 17). Querying this graph for a route between Busan and Singapore confines the search to a narrow corridor of feasible paths—indeed, in our case, Dijkstra's algorithm needed to explore only 32 nodes for the least-CO_2 solution (and 19 nodes for a purely shortest-distance path)—so the computation runs in essentially real time. Figure 1 provides a map overview of the extracted network, with nodes located at major turning points (such as approaches to the Taiwan Strait) and edges tracing along established shipping lanes.

3.2 Environmental Modelling and Route Optimization:

To evaluate route costs under specific weather conditions, the above network was coupled with forecast environmental data and a ship performance model. A representative late-spring (May) scenario was selected — corresponding to the early Southwest-monsoon onset that typically develops in May. During May, strengthening south-westerly winds associated with the monsoon transition often roughen the central South China Sea, whereas coastal waters closer to the Asian mainland can be relatively sheltered. No severe tropical cyclone is present, but there is a persistent gradient: rougher conditions to the north and milder conditions to the south, which is expected to influence optimal routing (ships may prefer more sheltered or southerly paths to avoid high waves).

A representative cargo ship was assumed for all routing computations to ensure consistency. Key particulars of

the vessel include a length of ~250 m, a design cruise speed of ~18 knots, and propulsion by a single heavy-fuel-oil engine. The ship's fuel consumption and speed behavior in various sea conditions were modeled with a simplified parametric approach. For baseline calm-water performance, standard empirical formulas were used (e.g., the Holtrop & Mennen (1982) method) to estimate the ship's resistance and power requirements at 18 knots. Weather impacts were then superimposed via heuristic adjustments to the effective speed on each leg. Head seas and strong headwinds induce added resistance that slows the vessel, whereas following seas, tailwinds, or favorable currents can increase its speed. For example, a head sea with significant wave height H_s causes a speed loss approximately proportional to H_s and the encounter angle:

$$\Delta V_w = k_w H_s \cos\theta \quad (1)$$

where θ is the relative wave heading (0° for head-on seas) and k_w is an empirical coefficient. Similarly, a headwind of speed U results in an added resistance corresponding to a speed reduction:

$$\Delta V_a = k_a U \cos\phi \quad (2)$$

With ϕ the relative wind angle and k_a a constant. (Following waves or tailwinds, which yield $\cos\theta < 0$ or $\cos\phi < 0$, produce a negative ΔV, i.e. a speed gain.) Any ambient current ΔV_c is directly added or subtracted along the ship's direction of travel. Thus, the effective transit velocity ΔV_{eff}, V_{eff} over a given segment is:

$$V_{eff} = V_{calm} - \Delta V_w - \Delta V_a + V_c \quad (3)$$

where V_{calm} is the vessel's calm-water speed (18 knots in this study). For each graph edge (segment) of length D, the transit time is then $T = D/V_{eff}$. Assuming the engine maintains constant power output corresponding to V_{calm}, the fuel consumption on that segment is taken proportional to transit time. If $\dot{m}_0$ is the fuel burn rate (tons per hour) is in calm conditions, then the fuel used over the segment is:

$$M_f = \dot{m}_0 T \quad (4)$$

This simple performance model is loosely based on standard industry methods for predicting added resistance in waves (ITTC, 2017). It captures the first-order effects of weather on speed and fuel use and was deemed sufficient for the present analysis. Given this model, each edge in the route network can be assigned an estimated transit time and fuel cost under the forecast weather conditions.

With the above components in place, route optimization is performed as follows. Each edge of the AIS network graph is weighted by either travel time or fuel consumption and CO_2 emissions for the given weather scenario. A standard shortest-path algorithm (Dijkstra's algorithm) is then run on the graph to find optimal routes according to different objectives. In this study, we consider two objectives: minimum fuel/CO_2, and (for completeness) minimum distance. The minimum solution on the high-resolution grid (without AIS constraints) is also evaluated as a baseline. The outputs of the routing algorithm are then analyzed for realism (overlap with actual traffic) and efficiency (transit time and fuel use).

3.3 Benchmark Routing Methods:

To evaluate the benefits of the AIS-constrained approach, we benchmarked it against two alternative routing methods that do not use the AIS network.

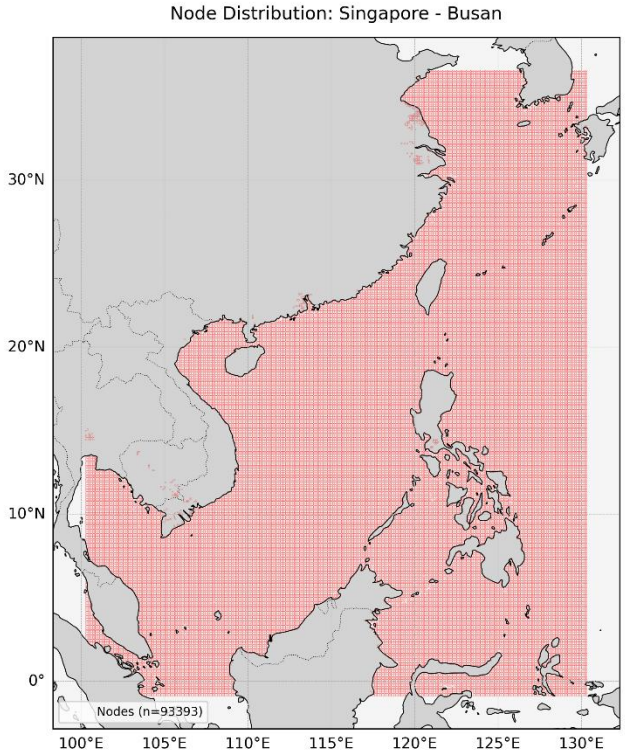

Figure 2. VISIR-2 regular grid (n = 93,393 nodes) for the Busan–Singapore benchmark

Grid-Based Weather Routing: First, we implemented a conventional grid-based routing algorithm. The reference grid-based solver employed a uniform 0.083° × 0.083°. In this approach, the ocean is discretized into a grid of waypoints (e.g. at regular latitude/longitude intervals). As

illustrated in Figure 2, the Busan–Singapore study domain is discretised with the original VISIR-2 0.083° × 0.083° grid (n = 93,393 nodes only for on the ocean), rather than with AIS-derived waypoints.

We performed a graph search over this grid to find the optimal path. For a fair comparison, we allowed the grid-based method to optimize for minimum travel time (fastest route) under the same environmental conditions. This roughly simulates the recommendation a traditional least-time weather routing system or experienced navigator might give, without considering fuel explicitly. The grid was fine-grained enough to capture route deviations for weather (on the order of 0.083° cells), and landmasses were treated as obstacles. The resulting fastest-route solution provides a baseline for minimum voyage time and a benchmark for fuel consumption if one prioritizes time over fuel. Figure 3 illustrates the three weather-optimised paths computed on the high-resolution grid (93 k nodes). This visual comparison highlights how each optimisation goal steers the vessel differently even within the same regular-grid framework.

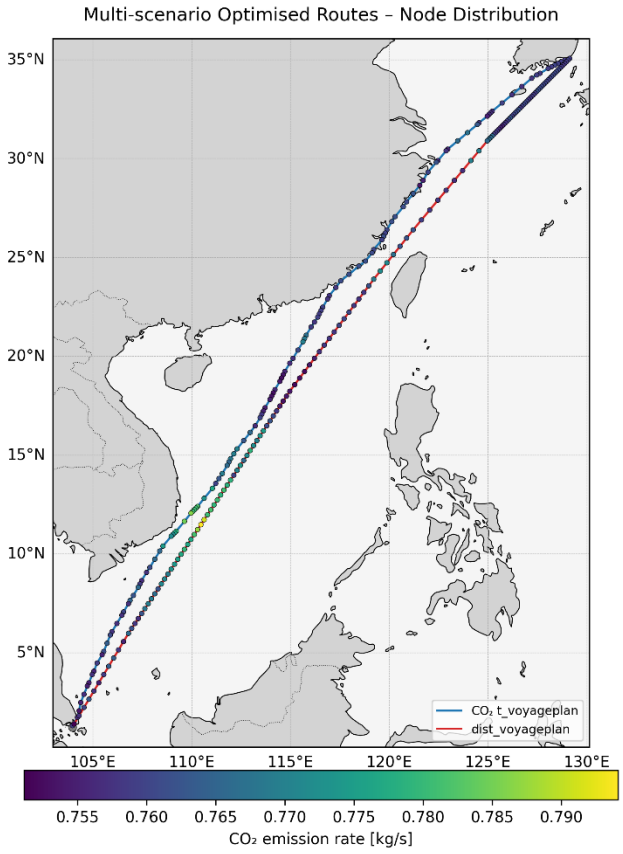

Figure 3. Multi-scenario weather-optimized routes on the Grid-Map

The regular VISIR-2 grid (93 393 nodes; 0.083° spacing) produces three distinct solutions (time-opt, distance-opt, CO_2-opt). Colour bar gives segment-wise specific-CO_2 [kg nm^{-1}].

VISIR-2 Optimal Routing: Second, we compared against an advanced weather routing model, VISIR-2. VISIR-2 is an open-source ship routing system that can compute paths optimizing various criteria (time, safety, etc.). We used the VISIR-2 model (Python implementation by Salinas et al., 2024) in a mode that minimizes fuel consumption and CO_2 emissions. This gave us a state-of-the-art unconstrained least-emission route for the scenario, against which we could measure our own results. The VISIR-2 model accounts for detailed wave-induced effects and was configured with the same vessel parameters and environmental inputs for consistency. We note that VISIR-2's solution is not restricted to historical lanes – it can choose any path across the water if it respects land and safety constraints.

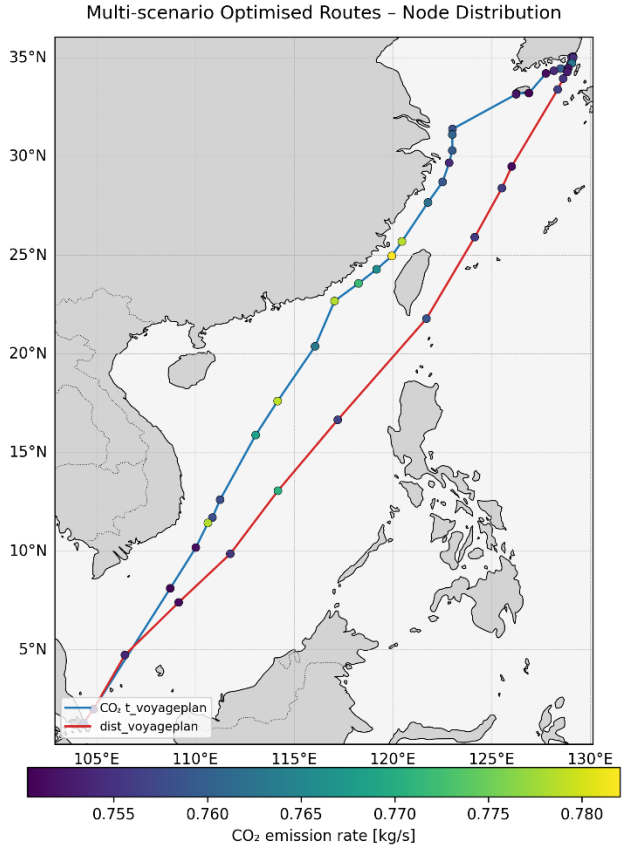

Figure 4. Multi-scenario weather-optimized routes on the MTN

The AIS-derived MTN collapses the search space to only 32 nodes (CO_2-opt) or 19 nodes (distance-opt) yet returns routes that hug real traffic corridors while matching, and in part surpassing, the grid solutions in fuel–time trade-off. Colour coding identical to Fig. 3 to allow direct comparison.

All three methods (AIS constrained, grid distance optimal, VISIR-2 fuel optimal) were applied to the same Busan–Singapore voyage scenario. The voyage was defined as a direct departure from Busan and arrival at Singapore, with no intermediate stops. We did not impose

any schedule requirements (no specific arrival time target); each method was free to extend the voyage if it yielded benefits in the optimized metric (e.g. fuel saving). This way, we could observe the pure trade-offs between time and fuel that each approach makes. All computations were done for the May transition-season scenario described, using identical weather forecast fields for fairness. Figure 1 shows the MTN extracted from AIS, which forms the foundation for our method. Figure 4 shows the corresponding paths on the MTN, where the optimiser explores barely 32 nodes yet still attains the same three objective functions.

4. Results and Analysis

4.1 Route Comparison

The routing outcomes for the scenario reveal clear differences in path selection between the methods (Figure 3, Figure 4). The fastest route (red baseline) computed on a conventional grid took a direct great-circle path, aiming to minimize transit time. This unconstrained route initially ignored some real-world traffic constraints: for instance, it did not strictly adhere to the established traffic separation scheme through the Taiwan Strait. In contrast, the AIS-constrained optimal route naturally followed the common maritime corridor through the Strait and along other high-traffic lanes. It chose a slightly more southerly course than the great circle at certain points, to avoid the worst wave conditions, but otherwise stuck to well-traveled sea lanes. The VISIR-2 fuel-optimal route was broadly like the AIS-guided route for much of the voyage, as that model also tended to avoid the rough northern waters. However, VISIR-2 detoured even further south in one portion of the route to squeeze out additional fuel savings, resulting in a more circuitous path. Consequently, VISIR-2's solution had the longest distance and time of the three, and it strayed slightly off the densest historical corridor in the South China Sea.

To quantify route realism, we compared each optimized path to actual AIS tracks from 2024. The AIS-constrained route achieved over 90% overlap with real vessel trajectories (measured via Symmetrized Segment-Path Distance, per Besse et al., 2015), indicating that it very closely matches the paths that ships take. The VISIR-2 route was also largely aligned with common traffic patterns (on the order of 85–90% overlap with AIS data), since it converged to a similar solution for most of the voyage. However, VISIR's extra detour in the southern sector placed it slightly outside the busiest shipping corridor, which lowered its overall overlap by a few percentage points. Meanwhile, the unconstrained grid-based route (distance-optimal) had a much lower similarity to real traffic. Before post-processing, that fastest route cut through areas where ships normally would not go (such as an off-limits section of the Taiwan Strait). In practice, a mariner or routing service would have to adjust this path to comply with navigational constraints, effectively bringing it closer to the AIS-guided solution. These observations underscore the value of introducing AIS-derived waypoints into the routing process: the optimized plan remains near-optimal in efficiency while inherently satisfying real-world route constraints.

4.2 Transit Time and Fuel Consumption

Table 1 contrasts the grid and MTN solutions, with all performance gaps (Δ-metrics) measured against the distance-optimal grid route. Constraining the search to the MTN dramatically reduces the computational search space (from 93 k grid nodes to 32 explored MTN nodes). The MTN CO_2-optimal route emits 1.9 % less CO_2 than the distance-optimal grid route, The reduction in emissions requires a longer passage of 6.5 hours (a 3.16% increase). The trade-off means that for every percentage point of extra transit time, the route achieves approximately 0.6 percentage points of emission saving. While this is a clear improvement over the grid-based CO_2 optimization which resulted in higher emissions, it represents a more moderate efficiency gain than previously estimated. This seemingly counter-intuitive outcome on the grid platform can be attributed to the complex interplay between different environmental factors. The CO_2 optimization algorithm may prioritize a path with calmer seas to reduce added resistance and lower the hourly fuel consumption rate. However, if this same path traverses an area with strong adverse currents, the vessel's speed-over-ground (SOG) is significantly reduced. Consequently, the substantial increase in total transit time can outweigh the marginal benefit of a lower consumption rate, leading to a net increase in total fuel burned and, therefore, higher overall CO_2 emissions for the voyage.

While a ~2% emissions reduction on a single voyage may seem modest, in absolute terms it equates to many tons of heavy fuel oil saved (and dozens of tons of CO_2

avoided). Across an entire fleet and year of operations, such savings would accumulate significantly lower fuel costs and emissions. Importantly, these gains were achieved without forcing the vessel onto any unnatural or unsafe track—on the contrary, the optimized route stays on pathways a mariner would consider normal and navigationally sound. The results clearly demonstrate the value of injecting historical traffic knowledge into the routing algorithm: the ship can reap fuel and emissions benefits and sail a route that aligns with real-world practices.

Performance-gap notation. Following the Δ-notation introduced in VISIR-2 (Mannarini et al., 2024) we measure the percentage or absolute gap between any candidate route and the fastest feasible one (time-opt).

$$\Delta\, CO_{2\ time-opt} = \frac{CO_2 - {CO_2}^{time-opt}}{{CO_2}^{time-opt}} \times 100\% \quad (5)$$

$$\Delta\, t_{time-opt} = t - t^{time-opt} \quad (6)$$

Here t and CO_2 are the transit time and voyage emissions of the route under evaluation, while $t^{time-opt}$ and ${CO_2}^{time-opt}$ are those of the grid-based distance-optimal benchmark. A negative Δ indicates an improvement. This mirrors the generic dQ formulation in VISIR-2 Eq. (23), which uses the least-distance route as an operationally relevant baseline.

4.3 Realism and Compliance Analysis

The above results are based on a single representative weather scenario. To ensure the findings are not specific to one set of conditions, the routing experiment was repeated under multiple different weather cases. Ten additional simulations were conducted, varying the seasonal timing and environmental inputs (e.g. different months with calmer or more severe weather) for the same Busan–Singapore voyage. The performance improvements observed with the AIS-constrained approach remained consistent across these trials. In every scenario, the AIS-guided route required less fuel than the unconstrained distance-optimal route. The fuel reduction achieved by the AIS route ranged from about 1.8% to 2.9% relative to baseline, with a mean of approximately 2.5% and a sample standard deviation of 0.4 percentage points. For comparison, the unconstrained VISIR-2 model, as documented in its own case studies, shows average savings typically ranging from 0.6% to 2.2% but can achieve far greater, double-digit savings (up to 49%) under specific weather conditions. These higher savings in VISIR-2 are generally associated with more significant detours and consistently longer travel times, highlighting the different optimization priorities of the two models. This consistency suggests that the benefits of the AIS-network-constrained routing approach are robust under a variety of weather patterns, not just the specific May example. Even when conditions changed, the hybrid method reliably produced near-optimal fuel savings while keeping the vessel on realistic routes.

Table 1: Fuel–time performance of two routing objectives on the AIS MTN for the May Busan–Singapore scenario. Δ-values are relative to the distance-optimal grid route.

Platform	Optimization	Travel-time T (h)	ΔT (h)	ΔT (%)	Total CO_2 (t)	ΔCO_2 (t)	ΔCO_2 (%)
Grid	Distance	205.9	0.0	0.0	566.15	0.0	0.0
Grid	CO_2	207.9	+2.0	+0.97	570.01	+3.86	+0.68
MTN	Distance	210.0	+4.1	+1.99	571.33	+5.18	+0.91
MTN	CO_2	212.4	+6.5	+3.16	555.39	-10.76	-1.9

5. Discussion and Future Work

The case study above illustrates how integrating historical route data can improve the practicality of optimized ship routes. The proposed framework, by constraining optimization to a realistic Maritime Traffic Network (MTN), produces voyage plans that are not only efficient but also credible to mariners. However, the current model has several limitations that open avenues for future research.

5.1 Limitations of the Current Model

A primary limitation is the model's dependence on the

quality and density of the underlying AIS data. As Onyango et al. (2022) have noted, in remote ocean regions with sparse traffic, the AIS-derived network may be incomplete or fail to capture the full range of viable routes. Consequently, strictly adhering to the MTN might preclude the truly optimal weather route, especially under unusual or extreme weather events where an "off-network" path would be safer or more efficient. Network abstraction is only as robust as the data used to build it, a point emphasized by Varlamis et al. (2021).

Second, the current MTN was derived without distinguishing between different vessel types or sizes, a common simplification in network generation as discussed by Huang et al. (2025). Vessel characteristics significantly influence maneuverability and route selection; for instance, as observed by Lee and Cho (2022), smaller vessels often have different operational patterns than larger ones, with the latter being restricted to deeper channels. Our unified network does not capture these vessel-specific nuances.

Third, as pointed out in the review process, the widespread adoption of this system could lead to traffic congestion. If multiple vessels are simultaneously directed toward the same weather-optimized corridor, it could create new bottlenecks and potentially compromise navigational safety, a critical consideration for high-traffic areas.

5.2 Future Work and Extensions

The identified limitations suggest several promising directions for future research.

Flexible and Dynamic Routing: To address the issue of data sparsity and exceptional weather, the routing framework could be made more flexible. This could involve allowing limited "off-network" deviations when a significantly better path exists outside the established MTN, or by dynamically expanding the network with new waypoints in low-traffic areas. A framework like MATNEC, proposed by Bläser et al. (2024), which uses an environment-adaptive approach to define network nodes, could serve as a valuable reference. This hybrid approach would balance adherence to common practice with the need for adaptability.

Vessel-Specific Networks: To account for different vessel classes, future work could, as suggested by Huang et al. (2025), focus on developing multiple, vessel-specific MTNs or, more efficiently, a single unified network with class-dependent edge costs (e.g., transit time, fuel use). This would allow for the generation of routes tailored to a ship's specific performance characteristics and constraints.

Multi-Objective and Real-Time Optimization: Voyage planning is often a multi-objective decision. The current framework can be extended to support this by assigning a composite cost to each edge. For example, as explored in similar contexts by Filipiak et al. (2020) using evolutionary algorithms, a multi-objective approach could generate a Pareto front of optimal routes, presenting the decision-maker with a set of trade-offs between the fastest, the greenest, and the most cost-effective paths.

Dynamic Load-Balancing for Congestion Avoidance: To mitigate the risk of traffic congestion, the framework could incorporate real-time traffic density into its cost function. By leveraging live AIS data streams, the system could, as envisioned by Huang et al. (2025), dynamically increase the cost of traversing congested edges, thereby encouraging traffic dispersion and enhancing overall network efficiency and safety. This real-time update capability is a crucial next step for operational systems.

Integration with High-Fidelity Models: While the environmental response model used here was sufficient for demonstration, integrating more complex, high-fidelity ship performance models is a clear path for improvement. Advanced models, like the VISIR service described by Mannarini et al. (2016), account for complex ship dynamics and can be used to calculate the cost of each network edge with greater accuracy. Since the underlying graph-search framework is modular, a principle emphasized in the review by Huang et al. (2025), improvements in the physics-based cost models can be seamlessly integrated to yield more precise and reliable results.

By addressing these limitations and exploring these extensions, data-informed routing can evolve into a powerful, adaptable, and widely trusted tool for greener and more efficient maritime transport.

6. Conclusion

This paper presented a novel ship-routing approach that combines classical weather routing algorithms with AIS-derived MTNs to improve route realism without sacrificing carbon-reduction performance. In a case study

of the busy Busan–Singapore corridor, the method leveraged one year of AIS observations to construct a data-driven route network comprising 10,956 waypoints and 17,561 edges. Optimal pathfinding on this network (using Dijkstra's algorithm) was then able to find fuel-efficient routes by exploring only a tiny fraction of the global solution space, keeping computation fast. The resulting AIS-constrained route closely matched actual vessel tracks and achieved a notable fuel saving of 1.9% compared to the conventional fastest route. This came at the cost of a 3.2% increase in transit time, representing an important trade-off between efficiency and operational realism. While these savings cannot reach the maximum potential of unconstrained, state-of-the-art models like VISIR-2, which can achieve savings of up to 49% in some scenarios by taking routes mariners might not typically use, our method's core strength lies in its high compliance with established sea lanes and mariner practice. These findings demonstrate that embedding historical navigational intelligence into the optimization process can produce voyage plans that are both efficient and credible to mariners. The ship essentially receives optimized advice that looks and feels like a conventional route. This approach thereby bridges the gap between mathematically optimal routing and real-world practicability. By steering ships along proven traffic patterns while exploiting weather opportunities, the method can facilitate greener shipping operations that mariners are willing to adopt. Further development and scaling of such data-informed routing techniques could play an important role in the maritime industry's efforts to reduce fuel consumption and emissions without compromising safety or reliability.

Acknowledgements

This work was supported by the Startup Growth Technology Development Program – Didimdol Project, 2024 Second-Half Call (Project No. RS-2024-00443443) funded by the Ministry of SMEs and Startups (MSS, Korea).

References

Besse, P., Guillouet, B., Loubes, J. M., & François, R. (2015). Review and perspective for distance based trajectory clustering. *arXiv preprint arXiv:1508.04904*.

Zhang, S.K., Shi, G.Y., Liu, Z.J., Zhao, Z.W. and Wu, Z.L. (2018), Data-driven based automatic maritime routing from massive AIS trajectories in the face of disparity, *Ocean Engineering*, Vol. 155, pp. 240-250.

Onyango, S. O., Owiredu, S. A., Kim, K. il, & Yoo, S. L. (2022). A Quasi-Intelligent Maritime Route Extraction from AIS Data. *Sensors*, Vol. 22, No. 22.

Lee, J. S. and Cho, I. S. (2022), Extracting the maritime traffic route in Korea based on probabilistic approach using automatic identification system big data, *Applied Sciences (Switzerland)*, Vol. 12, No. 2.

Varlamis, I., Kontopoulos, I., Tserpes, K., Etemad, M., Soares, A. and Matwin, S. (2021), Building navigation networks from multi-vessel trajectory data, *GeoInformatica*, Vol. 25, No. 1, pp. 69-97.

Filipiak, D., Węcel, K., Stróżyna, M., Michalak, M. and Abramowicz, W. (2020), Extracting maritime traffic networks from AIS data using evolutionary algorithm, *Business and Information Systems Engineering*, Vol. 62, No. 5, pp. 435-450.

Holtrop, J. and Mennen, G. G. J. (1982), An approximate power prediction method, *International Shipbuilding Progress*, Vol. 29, No. 335, pp. 166-170.

Mannarini, G., Salinas, M. L., Carelli, L., Petacco, N. and Orović, J. (2024), VISIR-2: ship weather routing in Python, *Geoscientific Model Development*, Vol. 17, No. 10, pp. 4355-4382.

Mannarini, G. and Carelli, L. (2019), VISIR-1.b: ocean surface gravity waves and currents for energy-efficient navigation, *Geoscientific Model Development*, Vol. 12, No. 8, pp. 3449-3480.

Mannarini, G., Pinardi, N., Coppini, G., Oddo, P. and Iafrati, A. (2016), VISIR-I: small vessels - least-time nautical routes using wave forecasts, *Geoscientific Model Development*, Vol. 9, No. 4, pp. 1597-1625.

Mannarini, G., Turrisi, G., D'Anca, A., Scalas, M., Pinardi, N., Coppini, G., Palermo, F., Carluccio, I., Scuro, M., Cretì, S., Lecci, R., Nassisi, P. and Tedesco, L. (2016), VISIR: technological infrastructure of an operational service for safe and efficient navigation in the Mediterranean Sea, *Natural Hazards and Earth System Sciences*, Vol. 16, No. 8, pp. 1791-1806.

IMO (International Maritime Organization)'s website: https://greenvoyage2050.imo.org/technology/weather-routing, last accessed in July 2025.

ITTC (International Towing Tank Conference) (2017), *Recommended procedures for added resistance and propulsion in waves*, 28th ITTC Seakeeping Committee Report.

Huang, L., Wan, C., Wen, Y., Song, R., and van Gelder, P. (2025), Generation and Application of Maritime Route

Networks: Overview and Future Research Directions, *IEEE Transactions on Intelligent Transportation Systems*, Vol. 26, No. 1, pp. 620-637.

Bläser, N., Magnussen, B. B., Fuentes, G., Lu, H. and Reinhardt, L. (2024), MATNEC: AIS data-driven environment-adaptive maritime traffic network construction for realistic route generation, *Transportation Research Part C: Emerging Technologies*, Vol. 169, 104853.

Received **12 June 2025**
1st Revised **04 July 2025**
Accepted **05 July 2025**