Unveiling individual and collective temporal patterns in the tanker shipping network

Kevin Teo[1], Naomi Arnold[1], Andrew Hone[2], Michael Coulon[3],
Martin Ireland[3], Mauricio Santillana[4], István Z. Kiss*[1,5]

[1] Network Science Institute, Northeastern University London, London, E1W 1LP, United Kingdom
[2] School of Engineering, Mathematics & Physics, University of Kent, Canterbury CT2 7NF, United Kingdom
[3] AlphaOcean.ai Ltd, Preston Park House, South Road, Brighton, East Sussex, United Kingdom, BN1 6SB
[4] Machine Intelligence Group for the Betterment of Health and the Environment, Network Science Institute, Northeastern University, Boston, MA, USA
[5] Department of Mathematics, Northeastern University, Boston, MA 02115, USA
February 10, 2026

Abstract

The global oil tanker shipping network emerges from individual ship and fleet decisions driven by economic, environmental, and operational efficiency factors. However, most existing analyses of shipping networks rely on static, time-aggregated representations, which overlook critical temporal patterns that connect individual vessel routing strategies with both operational efficiency and system-wide cargo flows. To address this gap, we introduce a dual-scale framework complementing sequential motif analysis—which captures recurring patterns in ships' regional visit sequences—with Dynamic Mode Decomposition (DMD) to extract temporal dynamics spanning individual trajectories to global flow patterns. Using movement data from oil tankers across four vessel classes, we demonstrate that vessels exhibiting diverse regional exploration patterns spend up to 50% more time carrying rather than seeking cargo, indicating greater economic and environmental efficiency. At the system level, DMD analysis of regional cargo flows reveals distinct seasonal cycles with amplitudes averaging 16% between peak and trough periods. Major importing regions (Europe, East Asia) show synchronous annual demand cycles, while export regions (Middle East, South America) exhibit anti-synchronous patterns. These multi-scale temporal dependencies, undetectable through conventional static analysis, reveal actionable performance differentials that enable data-driven routing strategies to optimize both economic efficiency and environmental sustainability.

The global maritime network handles over 80% of international trade by volume, with crude oil and petroleum products alone accounting for nearly 30% of this market [1, 2]. Despite its economic significance and substantial environmental impact [3–6], maritime shipping networks remain understudied compared to other transportation systems, particularly regarding their temporal operational dynamics [7]. This research gap is especially pronounced for oil tanker networks, which exhibit fundamentally different operational characteristics from the more extensively researched container shipping networks where fixed schedules and regular route structures have made traditional complex network analysis particularly effective for understanding system-level properties [8].

Current maritime network research has predominantly focused on static, time-aggregated analyses, with four main analytical themes: trade and connectivity; hubs and centrality; vulnerability and robustness; and communities and spatial structure [9–11]. These studies established that global shipping is dominated by hub-and-spoke structures, where major ports are connected to numerous smaller ports, while other important ports act as regional gateways controlling the flow of goods between different geographic areas [8, 12–16]. Moreover, such analyses have also been used to compare changes in network connectivity and regional communities at different time periods, often before and after significant geopolitical events [17–19]. These studies are commonly interpreted through the lens of economic and ecological impacts [15, 20–22].

Despite these insights, static approaches that aggregate temporal data obscure critical operational patterns. When networks are analyzed as snapshots, essential sequential information is lost: the timing and order of vessel movements between regions, and the dynamic variation in cargo flows over time. This information loss masks valuable operational insights, as vessels following identical aggregate movement patterns between the same ports may employ fundamentally different sequential strategies that significantly impact their efficiency and market positioning. These different temporal routing patterns create distinct operational outcomes that remain invisible in static networks.

These temporal dynamics are especially critical in oil tanker movements, which are operationally different compared to container liners. Unlike container ships operating on fixed schedules along established routes, tankers operate in charter-based markets where individual vessels compete for cargo, creating highly flexible and responsive movement patterns [8]. These competitive dynamics generate precisely the type of sequential, temporal dependencies that static aggregation eliminates—as vessels dynamically adjust their routing sequences based on market opportunities and cargo availability. Recent work on oil tanker networks has examined traffic density [23], hub-and-spoke structures [24], and influential port identification [14], yet these studies maintain the static network paradigm and do not capture the temporal operational patterns that may distinguish successful from unsuccessful competitive strategies.

To overcome these analytical limitations and uncover the hidden temporal structure, we demonstrate that sequential motif analysis (which identifies recurring patterns of regional visits) and Dynamic Mode Decomposition (DMD, a method for extracting temporal patterns from time series data) reveal previously undetected patterns spanning both individual vessel operations and global flow dynamics. This dual-scale approach links routing strategies to efficiency outcomes and cargo flows to seasonal cycles. Specifically, we show that ships exhibiting diverse regional exploration patterns—captured through sequential motifs—achieve significantly better laden-ballast ratios, a key efficiency measure closely related to the industry-standard Energy Efficiency Operational Indicator. In parallel, DMD applied to regional cargo flow time series reveals that different ship classes operate with distinct regional footprints and specialized seasonal trading patterns. We identify previously uncharacterized synchronous and anti-synchronous relationships between maritime regions that correlate with economic and climate cycles. While these analyses operate at different scales—individual vessel behavior versus collective fleet dynamics—they provide complementary insights into the multi-scale temporal structure of maritime operations.

Results

Data

We obtained proprietary data from our partner company AlphaOcean, covering voyages taken by 3026 medium to large crude-oil and petroleum tankers from 2016 January to 2020 February. This consists of 452 Panamax, 1141 Aframax, 610 Suezmax, and 823 Very Large Crude Carrier (VLCC), listed in order of their size in deadweight tonnage. The dataset contains a list of laden legs (single trips where the ship is loaded with cargo), detailing the origin/destination locations and their respective departure/arrival times, as well as an identifier of the ship that completed the leg. This enables the reconstruction of the entire historical journey of an individual ship, illustrated in fig. 1, typically consisting of alternating laden and ballast legs. The definition of ballast legs is in the Methods section; additional data processing details are provided in Supplementary Information section 1.1 to 1.3.

While the locations provided are accurate up to the port-level, our analysis condenses this to a *regional* level, grouping 1337 ports into 26 regions. This aligns with industry practices of pricing freight rate indices on a region-to-region basis and reflects the uncertainty that exact discharge ports are often not specified to operators until late in the voyage. In total, we recovered over 192,000 shipping legs (both laden and ballast), split amongst 33,509 Panamax, 94,840 Aframax, 34,401 Suezmax, and 29,418 VLCC legs.

Ship routing performance indicators

The tanker ships in our study typically do not travel on fixed routes, nor do they service particular countries or regions exclusively. Instead, they are free to compete for cargo across different markets and regions, with their movements driven by economic considerations such as travel times, fuel and staffing costs, potential earnings, and other market opportunities. This operational flexibility presents a unique challenge for tankers: how do we evaluate the routing performance of ships?

A key challenge in evaluating tanker routing performance lies in the inaccessibility of comprehensive voyage and vessel data. We propose an accessible metric using temporal data—the *laden-ballast ratio* (LBR),

$$LBR = \frac{\sum_{l \in L} \Delta t_l}{\sum_{l \in L} \Delta t_l + \sum_{b \in B} \Delta t_b}, \tag{1}$$

where Δt_l is the at-sea duration of laden leg $l \in L$ (adjusted for port waiting times; see Methods for details and robustness checks) and Δt_b is the duration of ballast leg $b \in B$, where L and B are the sets of laden and ballast legs, respectively, of a ship's trajectory. The ratio above simply measures the proportion of time spent laden at sea out of the total time at sea. The LBR is designed to capture the relative amount of time a ship spends performing useful work (transporting cargo) versus non-revenue generating operations (repositioning to load cargo). By focusing solely on temporal patterns of laden versus ballast operations, the LBR excludes ship-specific technical factors such as fuel consumption patterns or engine specifications. This design allows us to avoid scenarios where a vessel with efficient engines but poor routing outperforms another vessel with less efficient engines but superior routing strategy. This enables fleet-level comparisons where comprehensive operational data remains unavailable.

The stratification of our analysis by ship type (Panamax, Aframax, Suezmax, VLCC) controls for substantial inter-class variability in vessel characteristics and operational profiles. Within each class, multivariate analysis controls for vessel age, effectively isolating routing efficiency from vessel-specific factors. Furthermore, as demonstrated in Supplementary Information section 1.4, LBR captures fundamentally similar efficiency patterns to the International Maritime Organization's Energy Efficiency Operational Indicator (EEOI) [25], which measures CO_2 emissions per unit of cargo work. While EEOI remains an industry gold standard, its calculation requires proprietary fuel consumption and precise route distance data typically unavailable in fleet-level research datasets. Our LBR measure, though simplified, provides a theoretically grounded and empirically useful measure that reflects revenue-optimizing incentives, enabling large-scale analysis of routing efficiency patterns across global shipping networks.

The laden-ballast ratio distributions in fig. 2(top) expose significant performance disparities across each ship class, with the upper quartile (top 25%) displaying substantially higher values compared to the lower quartile (bottom 25%). Comparing median values between these groups reveals that the upper quartile achieve laden-ballast ratios approximately 50% higher than the lower quartile. This clear difference in time spent carrying versus seeking cargo indicates significant heterogeneity in operational strategies within identical ship classes. Such performance differentials

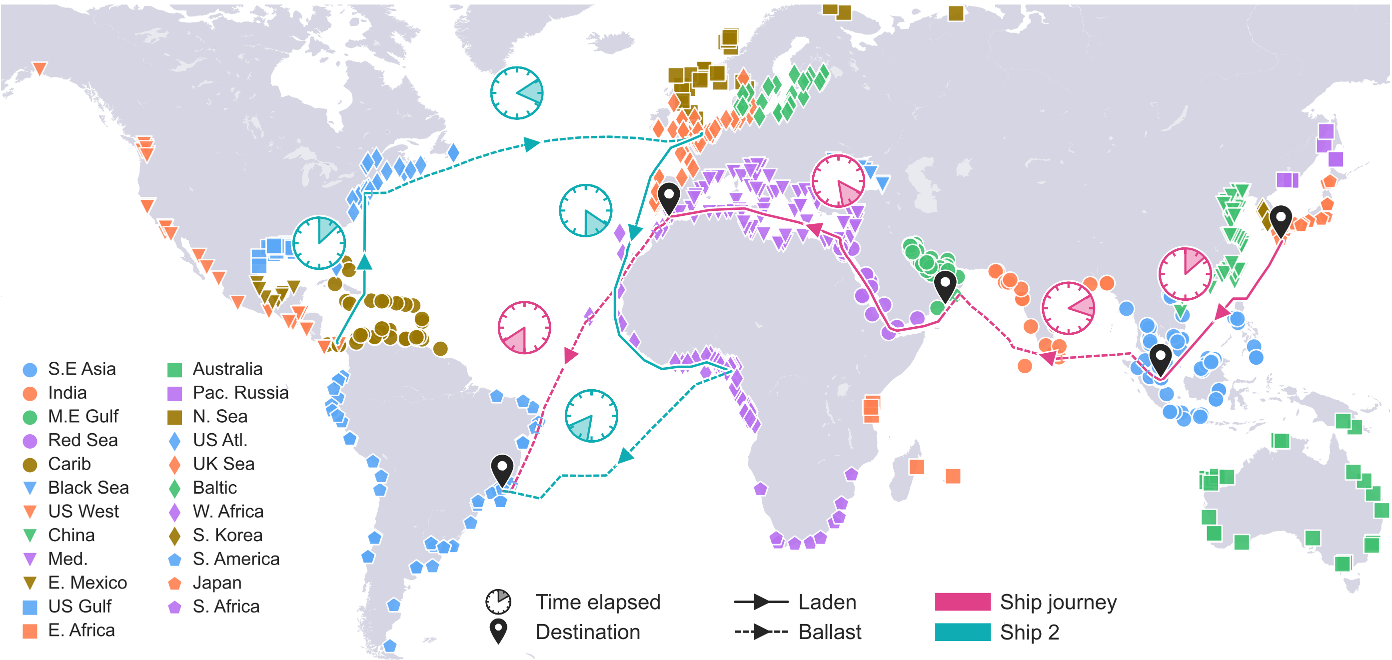

Figure 1: **Schematic representation of voyages in time and space of two ships.** Ships travel along directed port-to-port routes with journey time indicated by clock markers, alternating between laden and ballast voyages. Individual ports in the dataset are shown, with the regions they belong indicated by their marker shape and color.

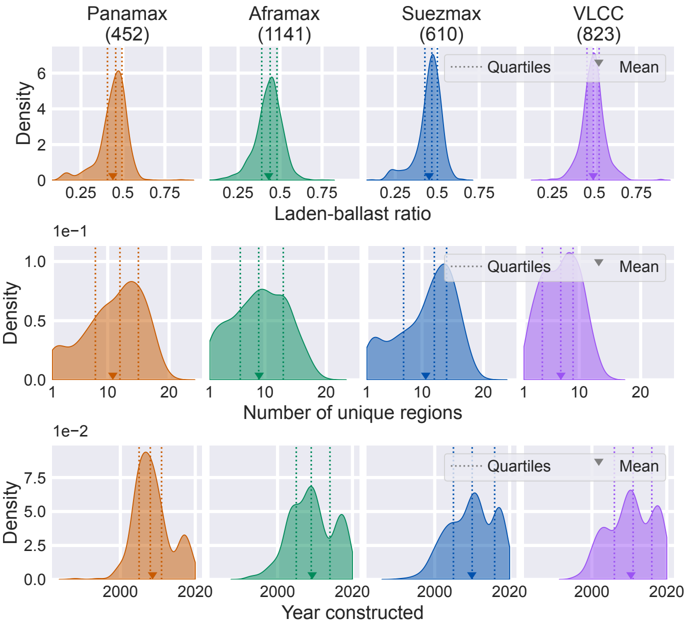

Figure 2: **Distribution of vessel measure data across classes.** Kernel density estimates for laden-ballast ratio (top), number of unique regions visited (middle), and year constructed (bottom) across four tanker classes. Dotted lines mark quartiles; triangles markers denote means. Sample sizes shown in parentheses: Panamax (452), Aframax (1141), Suezmax (610), VLCC (823).

suggest that optimizing routing strategies could substantially improve both economic outcomes and environmental performance.

This performance landscape reveals distinct efficiency profiles that align with each vessel class's operational niche. Panamax, Aframax, and Suezmax vessels show comparable laden-ballast ratios with overlapping median ranges, while VLCCs achieve marginally higher values due to their specialized operational advantages. As the largest vessel class, VLCCs are often deployed on inter-regional long-haul routes inaccessible to smaller ships, maximizing their at-sea laden time while minimizing time spent waiting in port. These longer voyage durations enable easier fixture planning, reducing inter-voyage delays that would otherwise erode efficiency.

Tanker vessels exhibit remarkable freedom to explore diverse markets rather than operating from fixed regional bases. Analysis of total regional coverage, fig. 2 (middle), uncovers distinct operational strategies across ship classes that directly reflect their design constraints and market positioning. Panamax and Suezmax vessels demonstrate similar exploration capabilities, visiting a median of 12 regions, indicating comparable operational flexibility despite size differences. Aframax vessels show the greatest variation in regional diversity with broader distribution patterns (median: 9 regions), reflecting their widespread usage, superior versatility for different cargo types, and port accessibility. VLCCs operate under significant geographical constraints, visiting only 7 regions on average with markedly narrow distribution. This restricted mobility stems from their specialized infrastructure requirements—as the largest vessels in the global fleet, VLCCs can only access ports with sufficient depth and handling capacity for their extreme size.

Extracting sequential motifs

Ship routing decisions involve complex factors that create distinctive movement signatures, exposing the strategic thinking behind high-performance operations. Consider two ships visiting identical regions: Ship 1 follows (*Middle East, China, Southeast Asia, Middle East, China, Southeast Asia*)—systematically cycling through three regions—while Ship 2 alternates (*Middle East, China, Middle East, Southeast Asia, China, Southeast Asia*)—return journeys between pairs of regions. In this example, time-aggregated networks would show identical connectivity between these three major trading regions, yet their sequential trajectories reveal fundamentally different strategies: systematic regional diversification versus back-and-forth movements between trading partners. To better understand the routing patterns that arise from these decision processes, we analyze short repeating sequences known as sequential motifs.

Previous studies have shown that ship trajectories can be modeled statistically using higher-order Markov chains, where future locations depend not only on the current position but also on recent prior locations [26, 27]. Higher-order networks generalize this correlation structure of individual agent trajectories to a network representation where nodes represent trajectory subsequences rather than individual ports, embedding these sequential dependencies in the network, where the order indicates the correlation length captured [28]. Utilizing this framework on our tanker dataset, we systematically identified order-2 as the optimal order based on information-theoretic model selection criteria (see Methods). Intuitively, this means that routing decisions depend on both the previous location and current location, but extending knowledge further back adds little predictive value while dramatically increasing model complexity. Building on the insights of this framework, we examine ship movements through 2-hop motifs to uncover the relationship between routing patterns and operational routing performance.

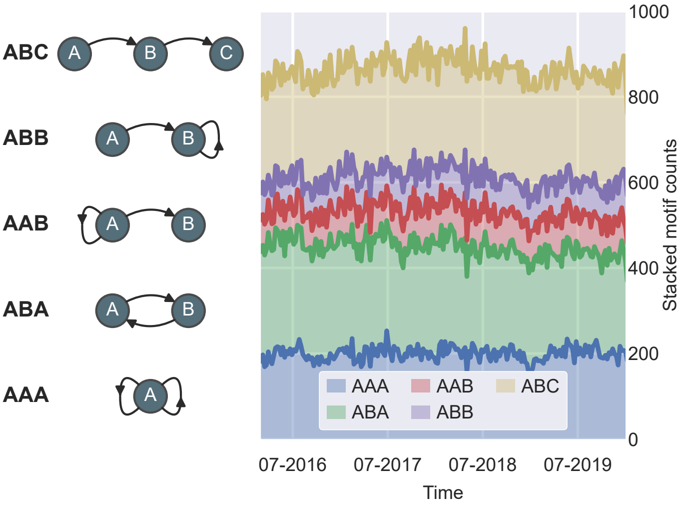

Figure 3: **Diagrammatic examples of 2-hop motifs and their observed counts over time.** (Left) Diagrams of 2-hop motifs. (Right) Stacked plot of motif counts over time in 1 week windows. Counts indicate the number of motifs that *started* within each 1 week window, to account for varying durations of 2-hop motifs.

Five distinct strategic patterns emerge from routing decisions. These are the five possible 2-hop motifs, which we call *AAA*, *ABA*, *AAB*, *ABB*, and *ABC*, illustrated in fig. 3 (left). The specific regional labels are not relevant; instead, motifs are identified by the topological relations between elements in the motif. For example, a journey sequence *(China, Japan, China)*, *(Japan, China, Japan)*, and *(UK, Mediterranean, UK)* all share the same underlying pattern: a ship visits region *A*, then region *B*, then returns to region *A*. This maps onto the motif *ABA*, regardless of the regions involved. Thus, *ABC* motifs represent maximally diverse routing (3 unique regions within 2 consecutive legs), while *AAA* motifs indicate routing focused within a single region.

We examine the absolute counts of different motifs over time, shown in the stacked time series in fig. 3 (right). Despite random fluctuations, autocorrelation analysis reveals no clear temporal patterns in any of the motif counts, suggesting that fleet-level routing behaviors remain stable and consistent over time (see Supplementary Information section 2.2 for autocorrelation plots). The motifs ranked based on occurrence frequency consistently reveal ABC as the most common motif, followed by ABA, and then AAA. This allows us to aggregate our motif counts across the whole duration of the data without losing critical information.

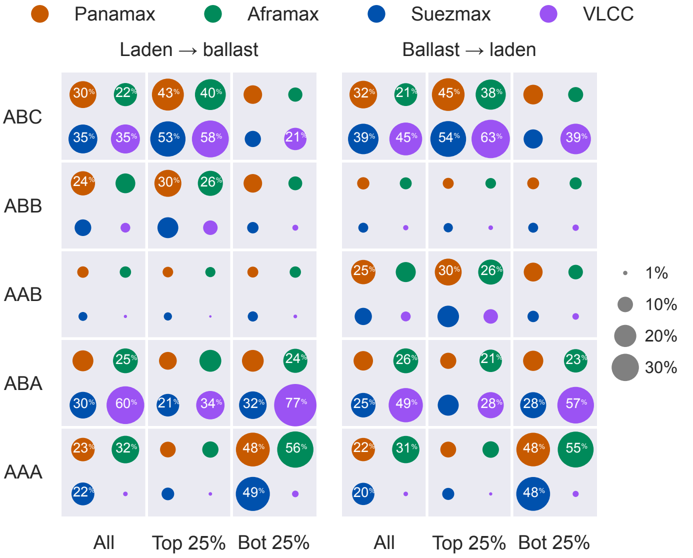

Figure 4: **Observed motif prevalence of each ship class based on laden-ballast ratio performance rank.** Motif prevalence by ship class for three groups: all ships, the top-ranked 25% by LBR performers, and the bottom-ranked 25% by LBR performers. Left panel shows laden-first motifs; right panel shows ballast-first motifs. Marker size indicates prevalence (in percentages).

To examine the relationship between routing patterns and LBR, we compared the motif prevalence between high-performing (top quartile by LBR) and low-performing (bottom quartile by LBR) ships, stratified by vessel class and cargo status (shown in fig. 4). Within each performance group, prevalence values are calculated as the weighted average of individual ship prevalence rates, where each ship is weighted by its total number of observed motifs; e.g. the weighted average prevalence of laden-first *ABC* motif among all Panamax ships is 30%.

Our analysis reveals a striking relationship: high-performing ships consistently exhibit greater prevalence of *ABC* motifs (visiting three different regions consecutively) and diminished prevalence of *AAA* motifs (remaining within the same region). Meanwhile, the inverse is true for low performing ships. This relationship holds across all ship classes, with top 25% performers displaying notably larger *ABC* prevalence and correspondingly smaller *AAA* prevalence compared to the bottom 25% performers (seen by the marker sizes in top 25% versus bot 25% columns).

The cargo status analysis reveals prominent usage of intra-regional legs for ballast operations. In laden→ballast motifs, *ABB* motifs (where the second leg represents intra-regional ballast movement) are more prevalent than *AAB* motifs overall, and especially in the top 25% group, consis-

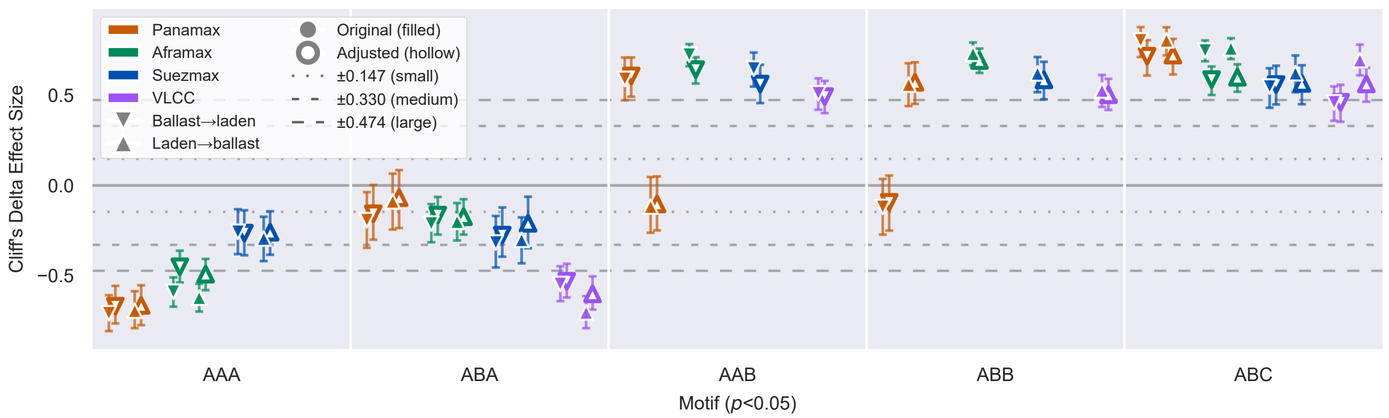

Figure 5: **Motif effect size (Cliff's delta δ) comparing top- and bottom- performers by laden-ballast ratio.** Effect sizes (with bootstrapped 95% confidence intervals) of statistically significant associations measured between top-ranked 25% against bottom-ranked 25%, grouped by vessel type and cargo status. Colors indicate vessel type and shapes distinguish cargo status (▲: laden→ballast; ▼: ballast→laden). Filled markers (▲, ▼) indicate original δ_{orig}, while hollow markers (△, ▽) indicate age-adjusted δ_{adj}. Dotted vertical lines denote effect size thresholds (small: ±0.147; medium: ±0.33; large: ±0.474).

tent with the strategy of minimizing ballast travel times by seeking nearby cargo when available. The reverse applies to the ballast→laden motifs, with increased *AAB* prevalence for high performance ships optimizing their cargo pickup location.

We strengthen our findings through comprehensive statistical validation. We systematically control for ship age—distribution shown in fig. 2 (bottom)—as a potential confounder and establish the statistical significance of observed trends through multiple approaches. Our stratified analysis by vessel class and motif cargo status employs Beta regression with motif prevalence as the dependent variable and LBR and ship age as regressors, isolating the independent contribution of each predictor (detailed in Methods and Supplementary Information section 2.3). The regression coefficients quantify each regressor's unique effect on motif prevalence, with p-values confirming whether these effects are statistically significant. Additionally, we capture the magnitude of group separation using Cliff's delta δ [29, 30], a robust non-parametric effect size that reveals distributional shifts beyond what simple mean comparisons can detect.

The associations between movement patterns and vessel LBR discussed above are statistically significant, show substantial effect sizes, and remain robust after controlling for vessel age. Effect sizes measured by Cliff's delta, shown in fig. 5, reveal large positive effects for *ABC* motifs across all vessel classes, while *AAA* motifs display small to large negative effect sizes for Panamax, Aframax, and Suezmax vessels. The laden-first *ABB* and ballast-first *AAB* motifs also exhibit large effect sizes across all vessel classes. These patterns achieve statistical significance ($p < 0.05$ after controlling for false discovery rates [31]) with the sole exception of *AAA* motifs in VLCCs. Critically, adjusting for ship age produces negligible changes in effect sizes (mean absolute difference $|\delta_{orig} - \delta_{adj}| = 0.056$; 20 out of 32 comparisons differ by less than 0.05), confirming that the motif-LBR associations represent genuine operational patterns rather than age-related artifacts. Qualitative interpretations of effect magnitude remain consistent whether or not age is explicitly controlled, with only Aframax's *AAA* ballast→laden shifting from large to medium effect size after adjustment (fig. 5).

VLCCs present a slightly unique case due to their operational constraint: unlike other ship classes, VLCCs show minimal *AAA* usage across both performance groups, reflecting their specialized role in long-haul routes with few intra-regional routes available to them. Instead, high-performing VLCCs show a reduced *ABA* motif prevalence (representing round-trips between two regions), suggesting that even within their operational constraints, high-performing VLCCs benefit from route diversification rather than repeating the same trading route. While reduced *ABA* motif prevalence of the top 25% is also observed in other ship classes, their effect sizes are notably smaller than in VLCCs, suggesting a weaker reliance on *ABA* motifs across the general population of ships (see Supplementary Information section 2.4 for VLCC leg duration comparisons of *ABC* and *ABA* motifs).

Together, these findings indicate that route diversification represents a key component of operational efficiency, with more efficient ships consistently exhibiting varied market exposure while optimizing ballast destinations through intra-regional movements.

Identifying periodicity in shipping activity

While our motif analysis establishes a connection between individual ship efficiency and route diversification, these individual routing decisions aggregate to create system-level phenomena. The collective movements of many ships, driven by opportunities arising from fluctuations in the crude oil market, manifest as non-trivial temporal signals at the regional level. Although some fluctuations are a result of highly unpredictable events (i.e. geopolitical disruptions or accidents), others may be driven by predictable physical or social phenomena, ranging from climate patterns [32] to socioeconomic patterns in production, operation, or consumption [33, 34]. We therefore shift our focus from ship-centric trajectories to examine the macro-scale cargo flow dynamics that these aggregated movements generate across regions.

These macro-scale dynamics are captured by time-series $y(t)$ of cargo flows (measured in $10,000$ deadweight tonnage) entering and leaving a region for each ship class, yielding 8 time series per region. Each unique combination [region, direction, ship class] is referred to as a *regional flow segment*, indexed with subscript n. To capture the complex dynamics within this collective time series $\mathbf{y}(t) = [y_1(t), \ldots, y_n(t), \ldots]$, we employ bagging-optimized Dynamic Mode Decomposition (DMD) [35–37], a powerful and noise-robust tool for extracting spatio-temporal patterns from a vector time series, with regional flow segments serving as spatial coordinates. Detailed methodology on the time series construction and DMD reconstructions is in the Methods section.

A rank R DMD yields R complex spatial components $\phi_{n,j}$, scalar amplitudes A_j, and a temporal frequency component ω_j, where n is the regional flow segments index and j is the index of the R modes. The DMD reconstructed time series for region flow segment n can be expressed in the form

$$\hat{y}_n(t) = \sum_j A_j |\phi_{n,j}| \cos\left(\omega_j t + \arg(\phi_{n,j})\right). \qquad (2)$$

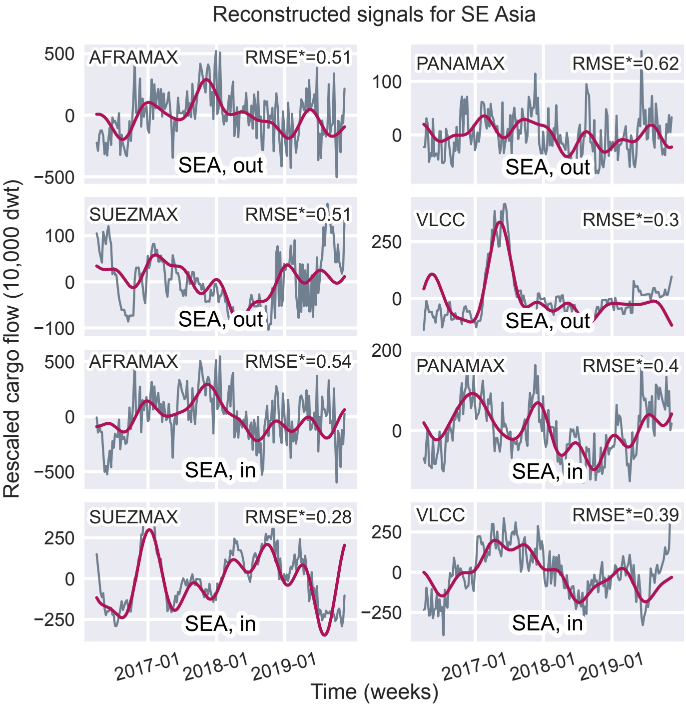

Figure 6: **Time series of cargo flow in South East Asia**. The original data is shown in gray, and the reconstructed signal in magenta. Each plot is labeled with the vessel class and the flow direction. The normalized root mean squared error $RMSE*$ is shown for each time series.

Despite the noisy data, the DMD method is able to consistently recover periodic oscillatory patterns; examples for the reconstruction $\hat{y}_n(t)$ for South East Asia are shown in fig. 6. The reconstruction errors are measured using the normalized root mean squared error $RMSE*$ [38], which lies in the range of $[0, 1]$ where 0 represents a perfect signal reconstruction. We see that for $RMSE* \leq 0.5$ the reconstructions visually fit the peaks and troughs, although much of the shorter-scale fluctuations are not captured.

The amplitude A_j is interpreted as the overall contribution of mode j to the dynamics; modes with larger A_j are more important. The spatial components $\phi_{n,j}$ represent the individual contributions of the regional flow segments n to mode j, and the complex argument of $\phi_{n,j}$ introduces a phase shift that controls when the peaks in each cycle occur. Thus, the combined term $A_j|\phi_{n,j}|$ in eq. (2) is interpreted as the amplitude of regional flow segment n. The temporal frequencies are more easily understood as cosine oscillation periods $\tau_j = 2\pi/\omega_j$. Note that, since DMD modes come in complex conjugate pairs with identical contributions to the reconstruction, we interpret our results in terms of these pairs.

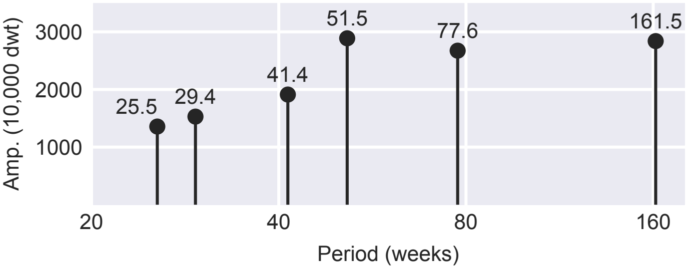

Figure 7: **Amplitude A_j against period τ_j from DMD decomposition**. Twelve modes were chosen for the decomposition, which were merged into 6 as each mode has a complex conjugate pair.

Our decomposition uncovers a dominant annual cycle that governs global tanker movements. The scalar amplitudes and oscillation periods from the decomposition are shown in fig. 7, where the mode index j are ordered in decreasing period lengths. The dominant mode-pair exhibits a striking 51.5-week oscillation period with 16% peak-to-trough amplitude variations across regional flows. To strengthen the validity of this finding, we applied two alternative periodicity extraction methods: Fourier transform and Multi-channel Singular Spectrum Analysis (MSSA) [39, 40]. Both methods independently recovered dominant periods of 48.0 weeks (95% CI: 47.1 – 49.9 weeks) via Fourier analysis and 49.5 weeks via MSSA, confirming the validity of our findings (see Supplementary Information section 3.2).

This yearly seasonal cycle firmly matches our prior expectations, as the natural seasons shape physical and socio-economic forces that drive the maritime oil trade. There are several direct factors that affect the movement of ships. In particular, the economic supply and demand of crude oil is driven primarily by the transportation sector, with secondary influences from the petrochemical industry and the energy sector [41, 42]. Analysis of global crude oil consumption [33] also confirms annual seasonality in consumption patterns. Our primary focus lies on interpreting this mode-pair; we present the rest our results in Supplementary Information section 3.3.

Our results reveal that different ship classes operate with distinct regional footprints, creating specialized seasonal trading patterns across the globe. Analysis of regional contributions to annual seasonality in each ship class, visualized in fig. 8, identifies where each vessel class concentrates its seasonal operations, with incoming and outgoing flow directions revealing the critical import and export hubs driving global trade cycles. The most significant patterns are described below.

For incoming flows, VLCCs exhibit the most regional concentration in seasonal amplitude variation, seen in fig. 8a, with China dominating the annual oscillations. This concentration reflects both the dedicated infrastruc-

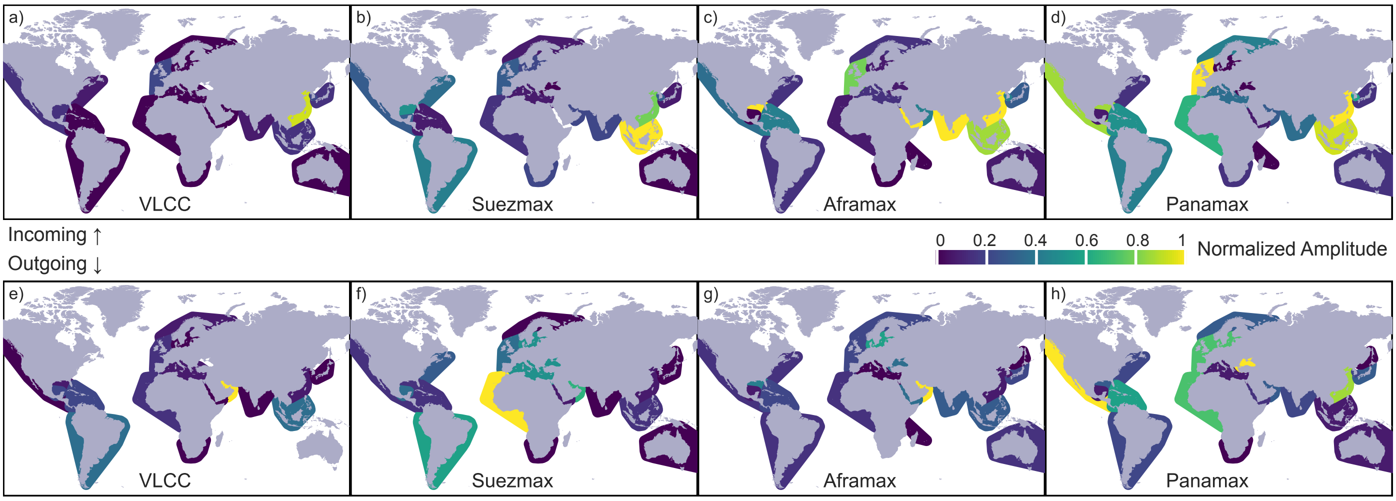

Figure 8: **Normalized regional amplitudes by ship class and flow direction.** Amplitudes, indicated by hue intensity, are normalized in each map by dividing actual amplitudes by the maximum observed amplitude within the same map. Maximum observed amplitudes by vessel class: VLCC (291.7 incoming, 238.7 outgoing), Suezmax (102.9 incoming, 99.0 outgoing), Aframax (99.6 incoming, 124.8 outgoing), and Panamax (38.2 incoming, 34.4 outgoing). Amplitude units are in $10,000$ deadweight tonnage.

ture required for the largest vessels and their substantial cargo capacity. These factors draw VLCCs to regions with the highest demand—primarily economically developed areas in the northern hemisphere with port facilities capable of accommodating these massive ships.

The geographical distribution of high incoming flow amplitude regions becomes increasingly dispersed for smaller vessels, reflecting their greater operational flexibility and broader port accessibility. Suezmax vessels (fig. 8b) maintain strong seasonal amplitudes around China but show their most prominent activity in Southeast Asia, with notable contributions from South America. Aframax vessels (fig. 8c) demonstrate further geographical spread encompassing much of Asia and central America, while Panamax vessels (fig. 8d) exhibit the most even distribution of amplitudes across the globe. The prominence of Western European and US ports for Panamax operations may reflect port size limitations in these regions. The amplitude scale factors (VLCC: 292, Suezmax: 103, Aframax: 100, Panamax: 38) reveal that a significant portion of seasonal demand is supplied by VLCCs, particularly toward China, while Aframax and Panamax vessels support most of the remaining global seasonal demand.

The outgoing flow patterns reveal complementary geographical specializations shaped by both resource distribution and maritime infrastructure constraints. VLCC exports (fig. 8e) show dominant seasonal amplitudes primarily in the Middle East followed by South America and Southeast Asia. While the Middle East and South America represent major crude oil production centers, Southeast Asia's prominence reflects its role as a major crude oil trading and storage hub rather than production, serving as redistribution centers for the Far East market. Suezmax (fig. 8f) exports are primarily driven by West Africa, the Mediterranean, and the Middle East—a pattern reflecting how Suezmax vessels, designed as the largest ships capable of Suez Canal transit, can efficiently serve seasonal Asian demand.

Aframax (fig. 8g) exports remain concentrated in the Middle East, while Panamax (fig. 8h) exports show peak amplitudes in the western United States, followed by Western Europe, West Africa, and China. The western United States dominance in Panamax exports suggests these vessels transited the Panama Canal eastward to access Atlantic regions [43]. These amplitude distributions demonstrate how size constraints interact with port infrastructure limitations and geographical resource distributions to create distinct seasonal flow patterns across the global shipping network.

Additionally, the DMD analysis on phase shifts uncovers patterns of synchronous and anti-synchronous oscillations between ship classes and between regions, illustrated in fig. 9. We find that few regions exhibit strong synchronicity across the ship classes, as observed through the clustering (or lack thereof) of markers along rows in fig. 9. Notably, only China has the phase shift of all ship classes being aligned around a seasonal peak in February/March, or $\pi/2$ corresponding to an offset of three months, (see fig. 9a).Other notable regions (Middle East and South East Asia) show moderate synchronicity, where only the VLCC outgoing flows display a distinctly different phase shift, highlighted in fig. 9b and fig. 9c.

We also observe clustering of multiple regions around the start of the calendar year, i.e. February/March (see fig. 9d), corresponding to the winter months of the Northern hemisphere. Indeed, these include prominent developed economic northern regions in Europe and Asia, which explains their ability to import in large quantities. Interestingly, South East Asia, while not a northern region, is largely in sync with this particular seasonal cycle, potentially due to their strong trading relationship with their northern Asian partners. Meanwhile, contributors to the outgoing flow (such as the Middle East) experience a peak around June/July or $-\pi/2$, anti-synchronous to the winter peaks. American seasonal patterns (USG, USW, and USA) do not appear to be in sync with the rest of the northern regions, potentially explained by their increasing independence in crude oil production [41].

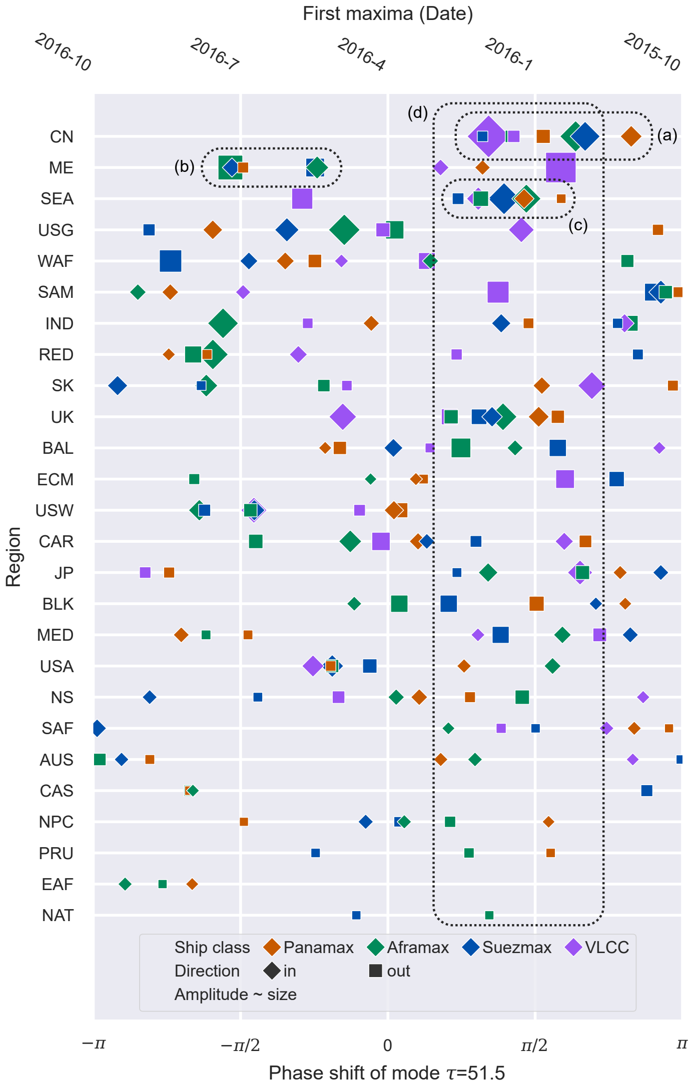

Figure 9: **Phase-amplitude plots for each region.** The color and marker size indicate the ship class and amplitude, respectively. The phase can be interpreted as where in time the first peak is by reading the top axis (note the direction of time is right to left), or the angular shift θ in $\cos(\theta + 2\pi/\tau)$ for $\tau = 51.5$ weeks by reading the bottom axis. Dotted outlines represent areas of interest.

Discussion

We demonstrated that moving beyond time-aggregated static network analysis reveals previously undetected patterns in maritime tanker networks at the individual vessel and fleet level. By applying sequential motif analysis, we uncover links between individual routing strategies and operational efficiency, while Dynamic Mode Decomposition shows seasonal spatio-temporal patterns in the dynamic global flow of cargo.

At the individual vessel level, our sequential motif analysis reveals that strategic route diversification correlates with superior operational routing performance across all ship classes. Top-performing ships favor *ABC* motifs (regional exploration) over *AAA* motifs (single-region focus) and *ABA* motifs (repeating round trips), even after controlling for ship age and stratifying by ship class, suggesting genuine strategic advantages beyond economies of scale or vessel-specific technical characteristics. While a definitive causal relationship warrants further investigation, we interpret this correlation as reflecting responsiveness to regional price differentials and cargo availability. Unlike container ships with fixed routes, tankers operate in charter-based markets where strategic flexibility enables exploitation of market opportunities across different regions, often manifesting as *ABC* motifs. Our findings suggest that fleet operators practicing regional specialization may improve performance through greater strategic diversification and market responsiveness. We note that *AAA* motifs encompass heterogeneous routing behaviors (i.e. from repetitive shuttles between specific port pairs to diverse movements across multiple ports within a region) which our regional aggregation cannot distinguish; future sub-regional analysis could clarify whether diversification benefits persist at finer spatial scales.

At the system level, our DMD analysis identifies pronounced seasonal patterns in regional cargo flows with quantifiable periods and phase relationships between maritime regions. These seasonalities represent strong underlying structural trends; however, they can also be obscured by short-term fluctuations week-to-week due to nonlinear drivers—including market volatility, port congestion, weather delays, or disruptive events—which manifest as high frequency noise. Our findings should therefore inform strategic planning over longer horizons. For fleet operators, this enables strategic positioning towards regions experiencing increasing seasonal supply or demand, opening opportunities for vessels willing to explore markets beyond their traditional operational areas. For port and governmental authorities, these patterns enable proactive planning for quarterly import/export variations, including timing infrastructure upgrades during off-peak periods, allocating seasonal workforce and berth capacity, and coordinating logistical resources to match anticipated cargo.

These findings establish strong ties between routing performance incentives and routing diversification; the preference for exploration over regional concentration among high-performing vessels provides actionable insights for fleet optimization strategies that can outperform competitors in revenue generation versus costs incurred. Furthermore, our new framework closes methodological gaps explicitly identified by Álvarez et al. [9], who emphasized the need for "deeper comparison among connected ports in terms of performance indicator" and urged researchers to "check the interactions between economy evolution in different regions and the maritime routes".

However, several limitations should be acknowledged. In the absence of ship-specific operational data, we constructed a performance indicator using the durations of laden and ballast legs. While these durations are recorded precisely in the data, they include port waiting times stemming from factors outside operator control (e.g., port congestion, administrative delays), which ideally would be isolated to strengthen inferences about routing strategy efficiency. With the availability of additional data—by including route-specific details such as the distance traveled and freight rate realized, ship-specific details such as the average speed and fuel consumption rate, or environmental conditions such as wind speeds or wave heights [44]—researchers could construct alternative performance indicators or EEOI proxies [45], including CO2-based indicators, enabling more direct evaluation of individual vessel routing decisions' environmental impact alongside economic efficiency. Alternatively, such enriched datasets would enable fleet-level comparative analysis through multivariate approaches that simultaneously control for technical, environmental, and operational factors while isolating routing strategy effects.

While the Bagging-Optimized DMD method incorporates sub-sampling to improve robustness against noise, our approach assumes an additive linear seasonal model [46], and deliberately excludes nonlinear drivers such as oil price shocks, policy changes, and geopolitical tensions or conflicts. Nonetheless, our primary objective was to establish whether underlying seasonal components exist and can be recovered despite system noise—and indeed we successfully demonstrate both. Future research could develop more predictive models that explicitly incorporate disruptive events alongside the baseline seasonal patterns we have established.

Our methodological choices, together with the pre-COVID time frame of our dataset (2016 January to 2020 February), limit the direct generalizability of the specific spatial and temporal patterns we report. Nevertheless, our dual-scale framework is fully deployable on post-COVID data, where comparisons to pre-COVID data can reveal whether major disruptions have fundamentally reshaped seasonal trading behavior or whether the system has reverted to its pre-disruption state.

Our methodology extends beyond oil tanker applications. Sequential motif analysis applies directly to dry bulk carriers and other charter-based shipping or transport sectors with flexible routing patterns. The DMD component could be applied to any transportation system with constrained capacity and regional flows, including trucking networks or airline cargo operations, while the general framework of linking individual behavioral patterns to system-wide efficiency outcomes has broad applicability across complex networked systems.

Methods

Ballast leg definition

Ballast legs (where a ship is not carrying cargo) are not explicitly recorded in our dataset; instead, we infer their occurrence between two consecutive laden legs of the same ship: given that a particular ship performs two laden legs from regions $a \to b$ and then $c \to d$, this implies a ballast leg from $b \to c$. Consequently, our definition of ballast legs may include other activities not typically related to reaching the next load port. Note that, as the only timestamps recorded are the departure and arrival times of each laden leg, the duration of ballast legs are consequently on average artificially longer than that of laden legs, as the additional amount of time spent on loading and discharging is not known. Nonetheless, including these port waiting times is logical from an economic perspective, as delays of this type incur further costs to operators without adding to revenues, much like ballast legs do.

Port waiting days adjustment

Our dataset records laden voyage departure and arrival times, but arrival timestamps correspond to berthing rather than port arrival, systematically conflating at-sea time with port waiting time. Since port waiting times vary randomly, this introduces randomly distributed positive bias that disproportionately affects shorter voyages.

We corrected this systematic bias using a reference dataset with complete temporal information (including separate port arrival and berthing times) covering different vessels and time periods. Analysis of this reference data revealed that port waiting times are approximately independent of voyage duration and appear randomly distributed within each vessel class. Empirically, these distributions are well-characterized by lognormal models, which we fitted using maximum likelihood estimation to obtain vessel-class-specific parameters (μ, σ) in log-space.

To ensure conservative corrections that avoid introducing bias in the opposite direction, we applied lower-bound estimates based on the mode of each lognormal distribution $\exp(\mu - \sigma^2)$. This approach identifies, with 95% confidence, the minimum port waiting time that vessels spend at the end of each laden leg. The conservative adjusted at-sea laden durations are obtained by subtracting the correction values from observed laden leg durations. See Supplementary Information section 1.3 for detailed results.

Validation analysis demonstrates that LBR performance rankings remain stable across correction levels: mean rank changes are ≤ 2 position for VLCCs, Suezmax, and Panamax, and ≤ 5 positions for Aframax vessels under median corrections (Figure S3 in Supplementary Information), with smaller changes for modal correction. For multi-event voyages (multiple loads or multiple discharges), we applied the same correction methodology at intermediate port visits to the total port duration from arrival to departure.

Comparison of laden-ballast ratio to EEOI

Performance indicators are useful tools for evaluating the efficiency of a ship's operations. The ideal performance indicator would accurately pick out the most efficient ships (e.g. high revenue yield, low emissions rate, low incurred costs, etc) and the least efficient ships accurately. The International Maritime Organization (IMO) encourages the use of the Energy Efficiency Operational Indicator (EEOI), which acts as an industry standard indicator. The EEOI is defined as

$$EEOI = \frac{\sum_i C_i}{\sum_i T_i D_i}, \tag{3}$$

where C is the carbon emissions, T is the cargo tonnage, D is the distance traveled, and i is the index of a leg (which can be either laden or ballast). The carbon emissions C are calculated as $\sum_\alpha F_\alpha c_\alpha$, where F_α is the quantity of fuel consumed and c_α is the CO_2 conversion factor for fuel type α. However, such detailed data are often only available to ship or fleet owners, making direct calculations and comparisons of the EEOI challenging on an industry level. Additionally, the EEOI is particularly sensitive to ship-specific parameters, and as such is not as useful in comparing ship-independent processes.

To tackle this challenge, we proposed an alternative route performance indicator, the laden-ballast ratio (LBR). In our data, the key parameters that are available to us are (i) the load departure time. and (ii) discharge arrival times. Given trajectory $J = \{(t_i, s_i)\}$ where t_i are port visit times and $s_i \in \{\text{load}, \text{discharge}\}$, we define the set of laden legs L and ballast legs B as

$$L = \{(t_i, t_{i+1}) : (s_i, s_{i+1}) = (\text{load, discharge})\}, \tag{4}$$
$$B = \{(t_i, t_{i+1}) : s_i = \text{discharge}, s_{i+1} = \text{load}\}. \tag{5}$$

Using these definitions, we compute the durations of the laden legs and ballast legs as

$$\text{For } l = (t_1, t_2) \in L\text{: } \Delta t_l = t_2 - t_1 - \mathrm{PWT}_c, \tag{6a}$$

$$\text{For } b = (t_1, t_2) \in B\text{: } \Delta t_b = t_2 - t_1, \tag{6b}$$

where PWT_c is the lower-bound estimated port waiting time for vessel class c. The LBR for a given ship is then defined as

$$LBR = \frac{\sum_{l \in L} \Delta t_l}{\sum_{l \in L} \Delta t_l + \sum_{b \in B} \Delta t_b}, \tag{7}$$

where Δt_l is the at-sea duration of laden leg l and Δt_b is the duration of ballast leg b.

The laden-ballast ratio measures the proportion of time a ship spends carrying cargo during its journey. This measure captures both the revenue-generating aspect of laden durations and the cost-savings associated with shorter ballast durations. By focusing on temporal patterns—time spent laden versus ballast—it isolates the effectiveness of routing decisions independent of vessel-specific factors such as engine efficiency or fuel-type. Ships with higher laden-ballast ratios spend more time performing useful work transporting cargo, or waste less time and energy traveling ballast, or both. As the industry increasingly adopts data-driven operations, this measure provides a practical tool for evaluating and comparing route-based performance between ships or fleets across different ownership groups.

Furthermore, the laden-ballast ratio shares similar properties with the EEOI. First, longer laden legs are beneficial. In the EEOI calculations, the average ratio of fuel consumed per unit distance for laden vs. ballast legs is much smaller than the ratio of tonnage carried (ballast legs have 0 useful tonnage carried). Thus, while laden legs emit more CO_2, longer laden legs tend be more efficient in the long run due to the tonnage carried being higher. Second, longer ballast legs are penalized. The ballast leg duration is defined as any time between consecutive laden leg durations. Thus, port waiting times are included in our ballast leg durations. As such, our performance indicator penalizes ballast legs more heavily than the EEOI—this approach reflects a financial perspective where time spent waiting at ports represents the same opportunity cost as time spent traveling ballast. See Supplementary Information section 1.4 for further details on the EEOI.

Sequential Motif Mining

Sequential motifs represent short recurring patterns within trajectory sequences that capture higher-order dependencies in movement data [47]. Unlike traditional network motifs that focus on structural relationships between ports [8, 13, 48–50], sequential motifs analyze the ordered trajectory patterns of individual ships navigating through the maritime network. Selecting an appropriate motif length (i.e. the chosen correlation length) is closely linked to identifying the optimal order in higher-order networks [26–28], where ℓ-hop sub-sequences serve as fundamental nodes of the network.

We determined optimal motif length by evaluating higher-order network models at orders 1 to 4 (beyond which model complexity becomes computationally expensive) using information-theoretic criteria that balance explanatory power against model complexity: Akaike Information Criterion (AIC) [51], Bayesian Information Criterion (BIC) [52], and likelihood ratio tests [53]. Each of these criteria applies different penalties, offering complementary perspectives on the optimal order selection. Parameters were estimated via maximum likelihood fitting of transition matrices on the full dataset for each vessel class.

Across all vessel classes, both AIC and likelihood ratio tests consistently selected order-2 as optimal (see Supplementary Information section 2.1). While BIC's stricter complexity penalty favored order-1, the order-2 BIC values remained substantially closer to order-1 than to order-3, indicating that order-2 provides meaningful improvement without excessive parameterization. Intuitively, this means that routing decisions depend on the previous location and current location, but extending knowledge further back adds little explanatory power at the cost of dramatically increasing model complexity. For order $\ell = 2$, the higher-order network edges represent 3-node trajectory subsequences, which our 2-hop motifs directly capture. This convergence of independent model selection criteria establishes that 2-hop motifs represent the minimal sequence length necessary to capture statistically significant sequential dependencies in tanker routing patterns.

Having established the optimal length, we now formally define sequential motifs. A sequential motif M of length ℓ is defined as a sequence of alphabets $(m_1, m_2, ..., m_{\ell+1})$ where $m_i \in \mathcal{A}$ for a finite alphabet set $\mathcal{A} = \{A, B, C, ...\}$. Denoting regions as v_i and the set of regions as V, a regional sequence $S = (v_1, v_2, ..., v_{\ell+1})$ constitutes an instance of motif M if there exists a unique mapping function $f : V \to \mathcal{A}$ such that $f(v_j) = m_j$ for $1 \leq j \leq \ell + 1$, with the constraint that $m_i = m_j \Rightarrow v_i = v_j$. This formulation ensures that motifs capture topological relationships independent of specific regional identities.

At length 2, there are exactly five possible motif types: *AAA*, *ABA*, *AAB*, *ABB*, and *ABC*. For example, the sequence (*China* → *Japan* → *China* → *SEA* → *SEA* → *Japan*) contained the following 2-hop motifs: *ABA*, *ABC*, *ABB*, *AAB*.

The total number of possible motifs of length ℓ follows the Bell numbers, representing unique partitions of $\ell + 1$ labels — see `https://oeis.org/A000110`.

Multivariate motif model

To control for potential confounding by vessel age, we employed multivariate Beta regression for each motif variant across all vessel classes. Beta regression is appropriate for modeling proportional data (e.g. motif prevalence $\in [0, 1]$) where observed values are close to the boundary [54, 55]. For each motif-class combination, we model motif prevalence ρ_M using Beta regression with logit link:

$$\mathrm{logit}(\mu) = \beta_0 + \beta_{LBR} \cdot \mathrm{LBR} + \beta_{age} \cdot \mathrm{age}, \tag{8}$$

where $\mu = \mathbb{E}(\rho_M | \mathrm{LBR}, \mathrm{age})$ is the conditional mean motif prevalence. The regression yields coefficients β_{LBR} and β_{age} quantifying the strength of association between each regressor and motif prevalence ρ_M. Statistical significance was assessed via hypothesis tests for $\beta_{LBR} \neq 0$ after applying Benjamini-Yekutieli correction across 40 tests (10 motif variants × 4 vessel classes) [31].

To isolate age effects from routing performance effects, we computed age-adjusted LBR residuals by regressing LBR against vessel age. These residuals capture the performance variation unexplained by age alone. Complete results on the

Beta regression coefficients, p-values, and age-adjusted effect sizes are detailed in Supplementary Information section 2.3.

Motif mining algorithm

For each vessel journey $S = (v_1, v_2, ..., v_{n+1})$ representing a sequence of regions visited:

1. Extract all consecutive 3-region subsequences $S_{i:i+2} = (v_i, v_{i+1}, v_{i+2})$ for $i = 1, ..., n-1$.
2. Apply the topological mapping function to classify each subsequence.

Letting $\kappa_M(S)$ = count of motif M in S, for a collection of journeys S_i, the prevalence of motif M is calculated as

$$\rho_M = \frac{\sum_i \kappa_M(S_i)}{\sum_i(|S_i| - 2)}. \tag{9}$$

Dynamic Mode Decomposition for shipping dynamics

To uncover coherent patterns in our regional shipping data, we employ the Dynamic Mode Decomposition (DMD) method [35], a powerful data-driven approach to studying dynamics of spatio-temporal data. DMD has been extended from its origins in fluid dynamics to a range of network data analysis—such as disease spreading dynamics and brain dynamics [56–58]—yet remains unexplored in maritime network analysis despite its suitability for spatiotemporal pattern extraction in constrained systems.

DMD is able to treat spatial and temporal modes simultaneously, making it particularly effective for identifying system-wide trends, something that can be leveraged in networks where the underlying topological connections act as a latent space for the spatial modes. We specifically use bagging-optimized DMD [36, 37]—implemented in `Python` with the `PyDMD` package [59, 60]—which improves robustness against noisy data through ensemble averaging. We provide a summary of the DMD output next; see [36, 37] for full details and theory.

Given a vectorized time series $\mathbf{Y} = \left[\mathbf{y}(1), \mathbf{y}(2), \cdots, \mathbf{y}(m)\right]$ over N regional flow segments and m time steps, the R rank DMD finds a complex-valued reconstruction

$$\hat{\mathbf{y}}(t) = \sum_{j=1}^{R} A_j \mathbf{\Phi}_j \exp(i\omega_j t) \tag{10}$$

where A_j is a scalar amplitude, $\mathbf{\Phi}_j = (\phi_{1,j}, \phi_{2,j}, \ldots, \phi_{N,j}) \in \mathbb{C}^N$ is a complex vector, and $\omega_j \in \mathbb{R}$ is the frequency of the j^{th} mode. The A_j, $\mathbf{\Phi}_j$, and ω_j terms are parameters to be fitted. In this case, the real part of $\hat{\mathbf{y}}$ is the best-fitted reconstruction of the data obtained by minimizing the Frobenius norm [37].

We can decompose the terms inside the right-hand side of eq. (10) into more intuitive forms. In particular, the real-component of the n^{th} region of the j^{th} mode can be written as

$$\Re[A_j \phi_{n,j} \exp(i\omega_j t)] = A_j |\phi_{n,j}| \cos\left(2\pi t/\tau_j + \theta_{n,j}\right), \tag{11}$$

where $\theta_{n,j}$ is obtained from the polar form of the complex number $\phi_{n,j} = |\phi_{n,j}| \exp(i\theta_{n,j})$, and $\tau_j = 2\pi/\omega_j$ is the period.

In order to apply DMD to our study, we must first get our shipping data in the form of $\mathbf{y}(t)$. Our data comes in the form of a set of laden legs. For a given vessel class, a laden leg describes a voyage from region u to region v between the times t_u and t_v with cargo weight w. This can be represented as a temporal edge (u, v, t_u, t_v). We then define a time step δt and rolling window interval ΔW, where $t_{i+1} = t_i + \delta t$. A temporal edge (u, v, t_u, t_v) is then considered active in interval $t + \Delta W$ if

$$\begin{aligned} t \le t_u &\le t + \Delta W \quad \text{edge begins in window} \\ t \le t_v &\le t + \Delta W \quad \text{edge ends in window} \\ t_u \le t \text{ and } t_v &\ge t + \Delta W \quad \text{edge overlaps window} \end{aligned} \tag{12}$$

Denoting the set of edges active in window $t + \Delta W$ as $\mathcal{E}(t, \Delta W)$, the time series of incoming or outgoing cargo flow for region i in time window t is the sum of the cargo weights w_e of all the active voyages $e \in \mathcal{E}(t, \Delta W)$:

$$\begin{aligned} y_{(i,in)}(t) &= \sum_{(u,v,t_u,t_v)=e\in\mathcal{E}(t,\Delta W)} \delta_{i,u} w_e \\ y_{(i,out)}(t) &= \sum_{(u,v,t_u,t_v)=e\in\mathcal{E}(t,\Delta W)} \delta_{i,v} w_e \end{aligned} \tag{13}$$

where $\delta_{i,u} = 1$ if $i = u$ else 0 and w_e is in units of $10,000$ deadweight tonnage (dwt). By definition then, ballast legs are not included. Thus, $y_{(i,in)}(t)$ is the total cargo arriving at region i at time t, and $y_{(i,out)}(t)$ is the total cargo departing from region i at time t.

An additional pre-processing step is to remove linear trends from $\mathbf{y}(t)$, as the DMD method can only recover sinusoidal and exponential terms. While the DMD algorithm may still converge without this step, it tends to fit eigenmodes with period $\gg$ the duration of the data set to approximate a linear term. This can obscure the functionality and interpretability of results, and thus it is better to remove it beforehand. We apply a linear regression in time to each time series, such that

$$\tilde{y}_i(t) = y_i(t) - r_i t - k_i \tag{14}$$

where r_i and k_i are gradient and constant terms from a least-squares straight line fit. This brings our work in line with an additive seasonal model [46], where the original signal $\mathbf{y}(\mathbf{t}) = \mathbf{L}(t) + \mathbf{I}(t) + \tilde{\mathbf{y}}(t)$, where $\mathbf{L}(t)$ is a long-term linear trend and $\mathbf{I}(t)$ are irregular terms. The component $\tilde{\mathbf{y}}(t)$ is composed of both the “seasonal” and “cyclical” terms, which in [46] corresponds to periodicities of < 1 year and > 1 year, respectively. However, in our case, we do not distinguish between the two.

Peak-to-trough relative amplitudes

To quantify the relative significance of an individual mode to each time-series, we need a standardized measure that accounts for the varying scales of cargo flows across different regional flow segments. To do this, we define the peak-to-peak relative amplitude, which normalizes each mode’s peak-to-peak amplitude by the total variability observed in the original time-series:

$$\text{Relative Amplitude}_{n,j} = \frac{2A_j|\phi_{n,j}|}{\max(y_n) - \min(y_n)} \tag{15}$$

where $A_j|\phi_{n,j}|$ represents the amplitude of mode j in regional flow segment n, and $\max(y_n) - \min(y_n)$ is the peak-to-peak range of the original time series for regional flow segment n. The factor of 2 converts the cosine amplitude to peak-to-peak amplitude. An additional factor of 2 is introduced when considering the relative amplitude of a mode-pair, since DMD modes come in complex conjugate pairs that both contribute to the reconstructed signal.

A larger value in the relative peak-to-peak amplitude indicates that the mode contributes more significantly to the observed fluctuations in that regional flow segment. The standardized nature of this measure also enables calculation of the average significance of each mode across all regional flow segments.

Choice of parameters

Dynamic Mode Decomposition requires the selection of several key parameters: SVD rank, time interval, rolling window duration, and number of delay embeddings. We employed a two-stage selection strategy utilizing error-based optimization (for SVD rank) and stability-based selection (for other temporal parameters) to avoid overfitting to data-specific noise.

For time interval (1 week), rolling window duration (6 weeks), and delay embeddings (22), we ensured that the parameter values chosen were statistically stable under partial variation. We systematically scanned parameter ranges while computing eigenvalue distributions for each combination, then calculated pairwise statistical distances (variational distance and Hellinger distance) between eigenvalue distributions across parameter sweeps [61, 62]. Regions of low statistical distance correspond to stable parameter regimes where extracted eigenvalues remain consistent across parameter perturbations (Figure S7 in Supplementary Information section 3.1). For time intervals and rolling window duration, we selected small parameter values to balance eigenvalue stability with maximizing the number of time points available for decomposition. For delay embeddings, we selected the smallest value within the stable regime that ensures the data matrix is tall (more rows than columns), which is required for stable SVD decomposition.

We determined the SVD rank through a two-step process. First, we examined the explained variance, identifying an "elbow point" where explained variance plateaus. Then we evaluated the root-mean-squared error (RMSE) around this elbow point, corresponding to candidate ranks $6 - 20$. The final rank was chosen such that it provided the largest average marginal improvement in RMSE across a range of delay-embedding. SVD rank 12 provided the largest average marginal improvement, while higher ranks showed inconsistent performance and diminishing returns. This approach balances reconstruction performance against model complexity to prevent overfitting to noisy signals.

The final parameter choices produced stable modes and eigenvalues that were well within the expected ranges under neighboring parameter choices. In particular, across all parameter combinations, we recovered the dominant annual mode with period 362 ± 10 days. Similar results were found for the other 5 modes. See our Supplementary Information section 3.1 for further details.

Reconstruction errors

The reconstruction errors for the $\tau = 51.5$ mode was evaluated for each combination of region, direction, and ship class using the normalized root mean squared error $RMSE*$ [38], shown in the scatterplot in fig. 10. The $RMSE*$ lies within the range of $[0, 1]$, with 0 indicating a perfect fit and 1 a maximally erroneous fit. We find that regions with larger amplitudes tend to have lower $RMSE*$, indicating that our interpretations should be focused on high amplitude regions.

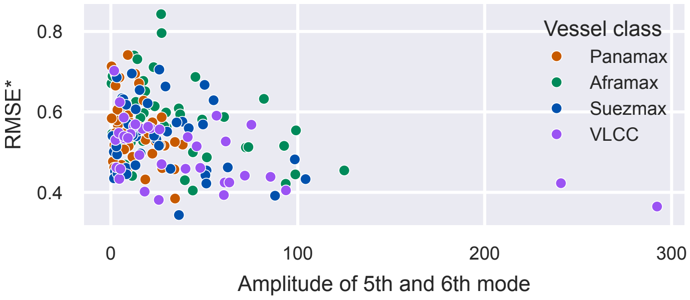

Figure 10: **Scatter-plot of reconstruction errors $RMSE*$ against 4th and 5th mode amplitude**. Lower values of $RMSE*$ indicate a better fit to the data.

Data availability

Due to legal restrictions, the data used to generate the main results in this article will be available only by contacting the author Kevin Teo.

References

1. United Nations Conference on Trade and Development. *Review of Maritime Transport 2017* 2017. `https://unctad.org/publication/review-maritime-transport-2017`.
2. U.S. Energy Information Administration. *World Oil Transit Chokepoints* 2024. `https://www.eia.gov/international/analysis/special-topics/World_Oil_Transit_Chokepoints`.
3. Rahim, M. M., Islam, M. T. & Kuruppu, S. Regulating global shipping corporations' accountability for reducing greenhouse gas emissions in the seas. *Marine Policy* **69,** 159–170. ISSN: 0308-597X. `https://www.sciencedirect.com/science/article/pii/S0308597X16301907` (2016).
4. Brancaccio, G., Kalouptsidi, M. & Papageorgiou, T. Geography, transportation, and endogenous trade costs. *Econometrica* **88,** 657–691 (2020).
5. Walker, T. R. *et al.* in *World seas: an environmental evaluation* (ed Sheppard, C.) Second, 505–530 (Academic Press, 2019). ISBN: 978-0-12-805052-1. `https://www.sciencedirect.com/science/article/pii/B9780128050521000309`.
6. Shi, Z. *et al.* Perspectives on shipping emissions and their impacts on the surface ocean and lower atmosphere: An environmental-social-economic dimension. *Elementa: Science of the Anthropocene* **11,** 00052. ISSN: 2325-1026. eprint: `https://online.ucpress.edu/elementa/article-pdf/11/1/00052/792716/elementa.2023.00052.pdf`. `https://doi.org/10.1525/elementa.2023.00052` (Oct. 2023).

7. Ducruet, C. The geography of maritime networks: A critical review. *Journal of Transport Geography* **88,** 102824. ISSN: 0966-6923. `https://www.sciencedirect.com/science/article/pii/S0966692320305895` (2020).
8. Kaluza, P., Kölzsch, A., Gastner, M. T. & Blasius, B. The complex network of global cargo ship movements. *Journal of the Royal Society Interface* **7,** 1093–1103 (2010).
9. Álvarez, N. G., Adenso-Díaz, B. & Calzada-Infante, L. Maritime Traffic as a Complex Network: a Systematic Review. *Networks and Spatial Economics* **21,** 387–417 (June 2021).
10. Faure, M.-A. & Ducruet, C. Shipping Network Research: A Systematic and Quantitative Review. working paper or preprint. `https://hal.science/hal-05106768` (Mar. 2025).
11. Ducruet, C. in *Port Systems in Global Competition* 300–333 (Routledge, 2023).
12. Hu, Y. & Zhu, D. Empirical analysis of the worldwide maritime transportation network. *Physica A: Statistical Mechanics and its Applications* **388,** 2061–2071. ISSN: 0378-4371. `https://www.sciencedirect.com/science/article/pii/S0378437108010273` (2009).
13. Ge, J., fu, Q., Zhang, Q. & Wan, Z. Regional operating patterns of world container shipping network: A perspective from motif identification. *Physica A: Statistical Mechanics and its Applications* **607,** 128171. ISSN: 0378-4371. `https://www.sciencedirect.com/science/article/pii/S0378437122007294` (2022).
14. Peng, P., Poon, J. P., Yang, Y., Lu, F. & Cheng, S. Global oil traffic network and diffusion of influence among ports using real time data. *Energy* **172,** 333–342 (2019).
15. Liu, Q., Yang, Y., Ke, L. & Ng, A. K. Structures of port connectivity, competition, and shipping networks in Europe. *Journal of Transport Geography* **102,** 103360. ISSN: 0966-6923. `https://www.sciencedirect.com/science/article/pii/S0966692322000837` (2022).
16. Zhang, Q., Pu, S., Luo, L., Liu, Z. & Xu, J. Revisiting important ports in container shipping networks: A structural hole-based approach. *Transport Policy* **126,** 239–248. ISSN: 0967-070X. `https://www.sciencedirect.com/science/article/pii/S0967070X22002086` (2022).
17. Ducruet, C. & Notteboom, T. The worldwide maritime network of container shipping: spatial structure and regional dynamics. *Global networks* **12,** 395–423. `https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1471-0374.2011.00355.x` (2012).
18. González Laxe, F., Jesus Freire Seoane, M. & Pais Montes, C. Maritime degree, centrality and vulnerability: port hierarchies and emerging areas in containerized transport (2008–2010). *Journal of Transport Geography* **24.** Special Section on Theoretical Perspectives on Climate Change Mitigation in Transport, 33–44. ISSN: 0966-6923. `https://www.sciencedirect.com/science/article/pii/S0966692312001561` (2012).
19. Guerrero, D., Letrouit, L. & Pais-Montes, C. The container transport system during Covid-19: An analysis through the prism of complex networks. *Transport Policy* **115,** 113–125. ISSN: 0967-070X. `https://www.sciencedirect.com/science/article/pii/S0967070X21003036` (2022).
20. Asgari, N., Farahani, R. Z. & Goh, M. Network design approach for hub ports-shipping companies competition and cooperation. *Transportation Research Part A: Policy and Practice* **48.** Psychology of Sustainable Travel Behavior, 1–18. ISSN: 0965-8564. `https://www.sciencedirect.com/science/article/pii/S0965856412001590` (2013).
21. Seebens, H., Schwartz, N., Schupp, P. J. & Blasius, B. Predicting the spread of marine species introduced by global shipping. *Proceedings of the National Academy of Sciences* **113,** 5646–5651. eprint: `https://www.pnas.org/doi/pdf/10.1073/pnas.1524427113`. `https://www.pnas.org/doi/abs/10.1073/pnas.1524427113` (2016).
22. Seebens, H., Gastner, M. T. & Blasius, B. The risk of marine bioinvasion caused by global shipping. *Ecology Letters* **16,** 782–790. eprint: `https://onlinelibrary.wiley.com/doi/pdf/10.1111/ele.12111`. `https://onlinelibrary.wiley.com/doi/abs/10.1111/ele.12111` (2013).
23. Yan, Z. *et al.* Analysis of global marine oil trade based on automatic identification system (AIS) data. *Journal of Transport Geography* **83,** 102637 (2020).
24. Peng, P., Yang, Y., Cheng, S., Lu, F. & Yuan, Z. Hub-and-spoke structure: Characterizing the global crude oil transport network with mass vessel trajectories. *Energy* **168,** 966–974 (2019).
25. International Maritime Organization. *Guidelines for Voluntary Use of the Ship Energy Efficiency Operational Indicator (EEOI)* Circular MEPC.1/Circ.684 (International Maritime Organization, London, Aug. 2009). `https://wwwcdn.imo.org/localresources/en/OurWork/Environment/Documents/Circ-684.pdf`.
26. Xu, J., Wickramarathne, T. L. & Chawla, N. V. Representing higher-order dependencies in networks. *Science Advances* **2,** e1600028. eprint: `https://www.science.org/doi/pdf/10.1126/sciadv.1600028`. `https://www.science.org/doi/abs/10.1126/sciadv.1600028` (2016).
27. Teo, K., Arnold, N., Hone, A. & Kiss, I. Z. Performance of higher-order networks in reconstructing sequential paths: from micro to macro scale. *Journal of Complex Networks* **13,** cnae050. ISSN: 2051-1329. eprint: `https://academic.oup.com/comnet/article-pdf/13/1/cnae050/61399620/cnae050.pdf`. `https://doi.org/10.1093/comnet/cnae050` (Jan. 2025).
28. Scholtes, I. *When is a network a network? Multi-order graphical model selection in pathways and temporal networks* in *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining* (2017), 1037–1046.

29. Cliff, N. Dominance statistics: Ordinal analyses to answer ordinal questions. *Psychological Bulletin* **114,** 494–509. `https://doi.org/10.1037/0033-2909.114.3.494` (1993).

30. Hess, M. R. & Kromrey, J. D. *Robust confidence intervals for effect sizes: A comparative study of Cohen'sd and Cliff's delta under non-normality and heterogeneous variances* in *annual meeting of the American Educational Research Association* **1** (2004).

31. Benjamini, Y. & Yekutieli, D. The control of the false discovery rate in multiple testing under dependency. *The Annals of Statistics* **29,** 1165–1188 (2001).

32. Stopa, J. E. & Cheung, K. F. Periodicity and patterns of ocean wind and wave climate. *Journal of Geophysical Research: Oceans* **119,** 5563–5584. `https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2013JC009729` (2014).

33. Inchauspe, J., Li, J. & Park, J. Seasonal patterns of global oil consumption: Implications for long term energy policy. *Journal of Policy Modeling* **42,** 536–556 (2020).

34. Raju, T. B., Chauhan, P., Tiwari, S. & Kashav, V. Seasonality in Freight Rates. *Journal of International Logistics and Trade* **18,** 149–157 (2020).

35. Schmid, P. J. Dynamic mode decomposition of numerical and experimental data. *Journal of fluid mechanics* **656,** 5–28 (2010).

36. Askham, T. & Kutz, J. N. Variable Projection Methods for an Optimized Dynamic Mode Decomposition. *SIAM Journal on Applied Dynamical Systems* **17,** 380–416. eprint: `https://doi.org/10.1137/M1124176`. `https://doi.org/10.1137/M1124176` (2018).

37. Sashidhar, D. & Kutz, J. N. Bagging, optimized dynamic mode decomposition for robust, stable forecasting with spatial and temporal uncertainty quantification. *Philosophical Transactions of the Royal Society A* **380,** 20210199 (2022).

38. Müller-Plath, G. & Lüdecke, H.-J. Normalized coefficients of prediction accuracy for comparative forecast verification and modeling. *Research in Statistics* **2,** 2317172 (2024).

39. Hassani, H. A brief introduction to singular spectrum analysis. *Optimal decisions in statistics and data analysis* (2010).

40. Hassani, H. & Mahmoudvand, R. Multivariate singular spectrum analysis: A general view and new vector forecasting approach. *International Journal of Energy and Statistics* **1,** 55–83 (2013).

41. U.S. Energy Information Administration. *Monthly Energy Review September 2024* tech. rep. Accessed: 2024-08-27 (U.S. Department of Energy, 2024). `https://www.eia.gov/totalenergy/data/monthly/`.

42. Wu, X. & Chen, G. Global overview of crude oil use: From source to sink through inter-regional trade. *Energy Policy* **128,** 476–486. ISSN: 0301-4215. `https://www.sciencedirect.com/science/article/pii/S0301421519300229` (2019).

43. U.S. Energy Information Administration. *Petroleum Supply Annual, Volume 1* tech. rep. Accessed: 2024-09-27 (U.S. Department of Energy, 2024). `https://www.eia.gov/petroleum/supply/annual/volume1/`.

44. Chen, X. *et al.* Intelligent ship route planning via an A* search model enhanced double-deep Q-network. *Ocean Engineering* **327,** 120956 (2025).

45. Chen, X. *et al.* Ship energy consumption analysis and carbon emission exploitation via spatial-temporal maritime data. *Applied Energy* **360,** 122886. ISSN: 0306-2619. `https://www.sciencedirect.com/science/article/pii/S0306261924002691` (2024).

46. Hylleberg, S. *Seasonality in regression* (Academic Press, 2014).

47. LaRock, T., Xu, M. & Eliassi-Rad, T. A path-based approach to analyzing the global liner shipping network. *EPJ Data Science* **11,** 18. `https://doi.org/10.1140/epjds/s13688-022-00331-z` (2022).

48. Si, R., Jia, P., Li, H. & Zhao, X. Assessing the structural resilience of the global crude oil maritime transportation network: A motif-based approach from network to ports. *Journal of Transport Geography* **123,** 104123. ISSN: 0966-6923. `https://www.sciencedirect.com/science/article/pii/S0966692325000146` (2025).

49. Xu, M., Deng, W., Zhu, Y. & LÜ, L. Assessing and improving the structural robustness of global liner shipping system: A motif-based network science approach. *Reliability Engineering & System Safety* **240,** 109576 (2023).

50. Wei, X. *et al. Resilience analysis of container port networks based on motif dynamics* in *2023 7th International Conference on Transportation Information and Safety (ICTIS)* (2023), 263–266.

51. Akaike, H. in *Selected papers of hirotugu akaike* 199–213 (Springer, 1998).

52. Neath, A. A. & Cavanaugh, J. E. The Bayesian information criterion: background, derivation, and applications. *Wiley Interdisciplinary Reviews: Computational Statistics* **4,** 199–203 (2012).

53. Wilks, S. S. The large-sample distribution of the likelihood ratio for testing composite hypotheses. *The annals of mathematical statistics* **9,** 60–62 (1938).

54. Ferrari, S. & Cribari-Neto, F. Beta Regression for Modelling Rates and Proportions. *Journal of Applied Statistics* **31,** 799–815. eprint: `https://doi.org/10.1080/0266476042000214501`. `https://doi.org/10.1080/0266476042000214501` (2004).

55. Geissinger, E. A., Khoo, C. L., Richmond, I. C., Faulkner, S. J. & Schneider, D. C. A case for beta regression in the natural sciences. *Ecosphere* **13,** e3940 (2022).

56. Proctor, J. L. & Eckhoff, P. A. Discovering dynamic patterns from infectious disease data using dynamic mode decomposition. *International Health* **7,** 139–145. ISSN: 1876-3413. eprint: `https://academic.oup.com/inthealth/article-pdf/7/2/139/16837591/ihv009.pdf`. `https://doi.org/10.1093/inthealth/ihv009` (Feb. 2015).

57. Griffith, T. D. & Hubbard Jr, J. E. System identification methods for dynamic models of brain activity. *Biomedical Signal Processing and Control* **68,** 102765 (2021).

58. Brunton, B. W., Johnson, L. A., Ojemann, J. G. & Kutz, J. N. Extracting spatial–temporal coherent patterns in large-scale neural recordings using dynamic mode decomposition. *Journal of Neuroscience Methods* **258,** 1–15. ISSN: 0165-0270. `https://www.sciencedirect.com/science/article/pii/S0165027015003829` (2016).

59. Demo, N., Tezzele, M. & Rozza, G. PyDMD: Python dynamic mode decomposition. *Journal of Open Source Software* **3,** 530 (2018).

60. Ichinaga, S. M. *et al.* PyDMD: A Python Package for Robust Dynamic Mode Decomposition. *Journal of Machine Learning Research* **25,** 1–9. `http://jmlr.org/papers/v25/24-0739.html` (2024).

61. Vaart, A. W. v. d. in *Asymptotic Statistics* 5–24 (Cambridge University Press, 1998).

62. Vaart, A. W. v. d. in *Asymptotic Statistics* 192–214 (Cambridge University Press, 1998).

Acknowledgments

Kevin Teo acknowledges the PhD studentship support from Northeastern University.

Author contributions

K.T, N.A, M.C, M.I, I.Z.K designed research; K.T, N.A performed research; M.S. contributed analytic tools; K.T, N.A, A.H, M.C, M.S, I.Z.K analyzed data; K.T, N.A, A.H, M.C, M.S, I.Z.K wrote the paper.

Additional information

Correspondence

Correspondence should be addressed to Kevin Teo or István Kiss. Materials request should be addressed to Kevin Teo.