Reliability Engineering and System Safety 279 (2027) 113821

Contents lists available at ScienceDirect

Reliability Engineering and System Safety

journal homepage: www.elsevier.com/locate/ress

A data-driven decision-making framework for dynamic pilot transfer point optimisation in port approaches

Gokhan Camliyurt [a], Bora San Turgut [b,c], Hyungoo Park [d], Sei Hun Kim [d], Daewon Kim [d], Youngsoo Park [d], Zaili Yang [a,*]

[a] Liverpool Logistics, Offshore and Marine Research Institute, Liverpool John Moores University, Liverpool, United Kingdom
[b] Transportation and Logistics Faculty, Istanbul University, Istanbul, 34452, Turkiye
[c] Department of Electrical Engineering, University of South Florida, Tampa, FL, 33620, USA
[d] Department of Navigation Sciences, National Korea Maritime and Ocean University, Busan, South Korea

ARTICLE INFO

Keywords:
Pilotage operations
Dynamic pilot transfer point optimisation
AIS trajectory analysis
Under-keel clearance (UKC)
Maritime traffic complexity

ABSTRACT

Maritime pilot transfer is among the most safety-critical manoeuvres in port approaches. This study presents a data-driven framework that predicts where a pilot transfer will take place when the pilot boat is dispatched, using only information available at that moment, and relocates the predicted point where hard safety constraints or a clearly better safety score require it. The framework integrates Automatic Identification System (AIS) trajectories, electronic navigational charts, tide-adjusted under-keel clearance (UKC), metocean observations and traffic context. Applied to one year at Busan Newport, it confirms 3243 transfers by the 13 port pilot boats against the official pilotage log. At a dispatch proxy identified retrospectively from AIS, inbound meeting distances are predicted with $R^2 = 0.495$ (mean absolute error, MAE = 1.28 km) under ship-grouped cross-validation and 0.537 in a forward temporal hold-out (0.30 and 0.32 without the transfer-time bound on the dispatch search); outbound distances are not predictable from the Busan AIS, which excludes the berths. The official fixed boarding zones perform far worse than any data-driven model. An out-of-fold optimisation using only AIS up to dispatch and chart-based routes finds a feasible recommendation in 96.4% of 1955 cases, all with UKC $\geq$ 2 m and 83.1% in the low collision-risk class, with a median relocation of 300 m. At 11 further stations in the United States and Denmark, locally retrained models reach inbound R^2 of 0.50–0.92 at ten stations. A size-corrected closest-point-of-approach (CPA) metric is checked against oriented hull polygons. The results support treating nominal boarding zones as dynamic corridors planned at pilot-boat dispatch.

1. Introduction

Pilot transfer remains one of the most safety-critical and workload-intensive phases of modern port navigation operations. During this manoeuvre, a vessel often needs to temporarily deviate from its routine trajectory to facilitate the safe transfer of a marine pilot under continuously changing environmental and traffic conditions. These operations frequently occur in confined waters with limited sea room, heterogeneous traffic flows, and complex bathymetric structures, which together heighten collision exposure during the approach phase. In practice, the transfer manoeuvre typically involves a deliberate alteration of course to provide a lee for the pilot boat, usually executed by turning opposite to the prevailing wind with a large rudder angle, briefly flattening the sea surface on the sheltered side and dampening pilot-boat roll and pitch. This calmness is, however, transient: the smoothened patch of water disappears within moments, and if the transfer cannot be completed within that short window the manoeuvre must be repeated. In severe weather the pilot may decline to disembark altogether, in which case the vessel proceeds to the next convenient port with the pilot still aboard. International regulatory frameworks such as International Maritime Organization (IMO) Resolution A.960 provide high-level criteria for determining safe pilot boarding positions, emphasising adequate distance from the pilotage initiation point, sufficient manoeuvring room, and avoidance of traffic convergence [1]. IMO Resolution A.1045 further specifies technical requirements for pilot ladders and transfer arrangements [2]. Crucially, the criteria contained in both instruments are expressed in verbal and qualitative form: terms such as “safe sea room,” “adequate manoeuvring room,” “safe from grounding,” and

* Corresponding author.
E-mail address: Z.Yang@ljmu.ac.uk (Z. Yang).

https://doi.org/10.1016/j.ress.2026.113821
Received 18 May 2026; Received in revised form 3 November 2026; Accepted 5 November 2026
Available online 6 November 2026
0951-8320/© 2026 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

"good/bad weather" are not accompanied by numeric thresholds and remain inherently subjective, leaving their interpretation to local pilot organisations, port authorities, and individual masters. This absence of quantitative limits introduces substantial inconsistency between ports and between operators, and is one of the gaps the present study seeks to close. However, both instruments implicitly assume that transfers occur at fixed or predesignated boarding locations. Such static spatial assumptions fail to reflect the reality of continuously varying sea state, visibility, under-keel clearance (UKC), or traffic complexity, which all together present strong influence to the operation safety. As a result, effective pilot meeting-point selection becomes a dynamic spatial decision problem influenced by environmental forcing, vessel interactions, and navigational constraints.

Professional guidance, including the International Maritime Pilots' Association (IMPA) and International Chamber of Shipping (ICS) Code of Safe Practice for Pilot Transfer Arrangements [3], and the operational principles described in the International Association of Marine Aids to Navigation and Lighthouse Authorities (IALA) Vessel Traffic Services (VTS) Manual [4], underscores the importance of communication integrity and situational assessment. Yet these materials similarly rely on qualitative judgement rather than quantitative, model-based decision support. Practical experience and accident investigations highlight the limitations of this approach. A widely cited incident in which a vessel struck a navigation beacon during pilot disembarkation, after situational awareness degraded while attempting to create a lee for the pilot boat, illustrates how quickly safety margins can collapse when environmental and traffic conditions change rapidly [5]. The Republic of the Marshall Islands Maritime Administrator has likewise documented recurring deficiencies near pilot stations, including over-reliance on VTS instructions, fatigue-related distraction, and failure to recognise no-go zones [6]. These incidents illustrate that pilot transfer safety is strongly influenced by dynamic environmental and traffic conditions, highlighting the need for quantitative decision-support tools capable of adapting pilot meeting locations in real time.

Furthermore, the operational context in port further exposes the discrepancy between formal guidance and real practice. For instance, although the port authority transitioned from a single fixed pilot boarding position to a broader polygonal boarding area to increase flexibility at Busan Newport, AIS analysis reveals that actual rendezvous locations are frequently situated outside these defined polygons, forming irregular and spatially dispersed patterns. Pilots and bridge teams routinely adjust their meeting points in response to prevailing sea state, hydrodynamic effects, and traffic density, particularly during heavy-weather conditions when pilot boats undergo violent rolling outside the breakwater. In such circumstances, transfers are often executed inside the breakwater despite regulations nominally defining external boarding positions. This raises a regulatory ambiguity: while compulsory pilotage boundaries specify where a pilot should embark or disembark, the framework rarely clarifies liability or compliance implications when operational necessity forces deviation. Whether a pilot boarding "too late," or disembarking "too early," constitutes a procedural breach remains a practical and legal grey area.

A further structural source of ambiguity is the command relationship itself. Under compulsory pilotage the master legally retains command of the vessel and the pilot is formally an advisor, yet in practice the pilot directs manoeuvres in real time on the basis of local knowledge while the master verifies and assents. This produces a de facto two-commander situation in which manoeuvring decisions are jointly authored, even though formal liability ultimately rests with the master and the vessel. As a result, when a transfer-related incident occurs, the subsequent inquiry is primarily a ship-accident investigation concerned with course, speed, and manoeuvre choice, rather than an occupational-injury investigation, although both dimensions may coexist and are treated separately in the liability framework introduced below.

When a pilot transfer accident does occur, post-incident inquiry typically proceeds in a structured order: (i) was the vessel under compulsory pilotage; (ii) did the transfer occur at, or sufficiently close to, the designated boarding location; and, if not, (iii) were there operationally defensible reasons for the deviation. The relevant risks can be grouped into five recurring categories: *grounding* (insufficient under-keel clearance); *collision* with other vessels in the encounter set; *allision* with fixed port infrastructure (breakwater, beacon, channel edge); *occupational injury*, in particular pilot falls during ladder transfer under unstable sea state or wind; and situational-awareness failure under reduced visibility. A framework whose variables map onto these categories can provide a structured, reproducible record of the conditions at a proposed transfer point, which may inform, but does not replace, operational judgement and post-incident inquiry. We return to this mapping in Section 6.1.

From the shipmaster's perspective, decision-making pressure begins long before the physical transfer. Even though AIS enables early identification of the pilot boat, effective rendezvous still depends heavily on visual acquisition, particularly at night or in complex approaches where background illumination or neighbouring traffic complicates lookout procedures. Bridge teams often begin visually searching for the pilot boat tens of minutes before the estimated meeting point, a practice rooted in operational necessity rather than regulatory prescription. This anticipatory search period adds cognitive load during a phase when navigational demands are already intensifying, especially under high traffic density or restricted visibility. Simultaneously, the transition from open-sea navigation to the boarding zone results in a marked increase in both bridge-team and deck-team workload: while open-sea navigation may be maintained by a single officer of the watch and one able seafarer, the approach to the pilot station typically requires the full navigational team, alongside deck personnel preparing the ladder, clearing moorings, and readying transfer equipment. Under adverse environmental conditions these preparations cannot be staged early, further concentrating human-factor stress during the most demanding navigation window.

Data-driven methods are increasingly applied across transport systems, for example generative artificial intelligence in intelligent transportation systems [7], predictive control of electric-vehicle charging demand [8] and traffic-parameter extraction from camera video [9]. Dynamic rendezvous and trajectory planning are established in air and land transport but do not address maritime constraints such as bathymetry, UKC, manoeuvrability and VTS-managed traffic, while maritime research on pilot transfer has focused on safety, ladder design, human factors and incidents [10,11] (Section 2).

This gap motivates the present study. Drawing on one full year of operational data from Busan Newport and 3243 pilot transfers confirmed against the official pilotage log, we develop a framework that predicts the meeting point from information available at pilot-boat dispatch and relocates it under hard safety constraints. The framework formalises the principles that experienced pilots and shipmasters implicitly apply: maximising sea room, maintaining UKC safety margins and avoiding traffic-induced complexity.

The study makes four contributions: (i) a decision-point formulation of meeting-point prediction at the dispatch of the pilot boat (identified retrospectively from AIS), with every input taken at or before that moment, validated against an official pilotage log and evaluated with ship-grouped and temporal partitions against fixed-zone and learned baselines at Busan and 11 further stations; (ii) a safety-constrained optimisation that anchors on out-of-fold predictions, routes the pilot boat and the ship through charted water, calibrates its gates within the training folds and relocates the predicted point when UKC, traffic or navigational constraints are violated, or reports that no feasible point exists; (iii) a size-corrected closest-point-of-approach metric, CPA-SD (size-corrected CPA), reported in signed and truncated form and checked against oriented hull polygons; and (iv) a mapping of the model variables to five categories of pilot-transfer liability. Together, these contributions shift pilot transfer point selection from a static chart-based practice toward a transparent, data-supported decision framework that

bridges static regulatory designations and the dynamic realities of pilot transfer.

The rest of the paper is organised as follows. Section 2 reviews the relevant literature on pilotage safety, AIS-based encounter analysis, hydrodynamic and bathymetric constraints, and dynamic meeting-point optimisation. Section 3 describes the data sources and preprocessing pipeline. Section 4 presents the methodology, including geometry-aware event reconstruction, the CPA-SD metric, and the hybrid Random Forest–Ridge optimisation framework. Section 5 reports the empirical results for Busan Newport. Section 6 discusses the findings, policy implications, and transferability to other ports. Section 7 concludes.

2. Literature review

Although pilot boarding grounds are typically designated as fixed charted locations, operational practice frequently requires vessels to adjust their rendezvous positions in response to prevailing sea state, traffic density, and manoeuvring constraints. Previous research has examined pilotage safety from several perspectives, including human-factor analysis, accident investigation, AIS-based maritime traffic analytics, and maritime encounter modelling. However, relatively limited attention has been given to the spatial and operational optimisation of pilot meeting points using integrated operational datasets, despite its curial role in navigation safety in port. One of the challenges lies in the high requirements on comprehensive supporting data involving ships, sea and weather conditions across a reasonable period time that can represent the complication of all different scenarios.

Early studies on pilotage safety primarily focused on human-factor analysis and accident causation mechanisms. Cognitive modelling approaches have been applied to maritime accident analysis, demonstrating the importance of situational awareness, decision workload, and organisational influences during critical navigation phases [10]. Empirical investigations of maritime incidents have likewise emphasised that bridge-team coordination, communication quality, and operational workload significantly affect navigational safety in complex environments such as port approaches [12]. More recent research applying machine-learning techniques to maritime safety has shown that environmental conditions, vessel characteristics, and operational context can influence the probability of pilot-related incidents and occupational accidents [13]. Studies focused specifically on pilot-transfer occupational risk have quantified the elevated likelihood of pilot ladder failures, falls, and crew-side procedural deficiencies through fault-tree and event-tree analyses [14], while structured human-error frameworks such as HFACS-PV combined with the Success Likelihood Index Method (SLIM) have been applied to the pilot-transfer process to identify the dominant individual and organisational error modes [15]. Bayesian-network and fuzzy fault-tree formulations have similarly been used to dynamically characterise pilot-transfer accident causation under varying operational conditions [16]. Complementary studies have examined the systemic influences on pilot decision-making [17] and the broader role of human factors in pilotage safety [18]. While these studies provide valuable insight into the human, organisational, and probabilistic dimensions of pilotage risk, they primarily focus on accident causation and decision context rather than on predicting or optimising the spatial location at which pilot transfers occur.

Accident investigations, introduced above [5,6], also link collisions and close-quarters situations near pilot stations to traffic congestion, restricted sea room and the limited manoeuvrability of large vessels [19,20]; they remain qualitative and give little guidance on where transfers should take place under varying conditions.

Parallel developments in maritime analytics have increasingly leveraged AIS data to analyse vessel movements and maritime traffic patterns. AIS-based studies have been widely used to detect near-miss collisions and encounter risks [21], estimate ship-interaction probabilities in dense traffic environments [22], and identify shipping routes or anomalous vessel behaviour through trajectory mining approaches [23]. AIS data have also been applied in environmental analyses, such as estimating maritime emissions based on vessel movement patterns [24] and their spatiotemporal distribution in port areas [25], and in machine-learning applications for classifying encounter situations and predicting vessel behaviour [26], including graph-based forecasting of ship speed in port [27]. More recent studies have further applied deep learning and advanced machine-learning techniques to AIS trajectory datasets to model vessel behaviour and navigational risk in complex maritime environments [28–30]. These studies demonstrate the growing value of AIS data for understanding vessel traffic dynamics and navigational risk. However, most AIS-based research focuses primarily on collision detection or traffic pattern analysis rather than modelling the operational decision processes that determine where pilot transfers occur within a port approach area.

Another important dimension of port navigation research concerns hydrodynamic and bathymetric constraints, particularly UKC. Studies on UKC modelling have demonstrated that tidal elevation, bathymetric uncertainty, vessel squat, and hydrodynamic interaction significantly influence navigational safety margins during port approaches and fairway transits [31]. These models are widely used in port engineering and channel design to ensure safe navigation for large vessels. However, existing UKC research primarily addresses grounding risk and fairway depth management rather than the spatial decision process associated with pilot boarding operations. Similarly, research on pilotage planning has examined procedural and operational aspects of pilot boarding arrangements [32], but these studies generally assume fixed pilot boarding zones and do not consider dynamic relocation of pilot meeting points in response to operational conditions.

Recent advances in machine learning and data-driven maritime analytics have created new opportunities for modelling complex navigational behaviour. Machine-learning approaches have been applied to accident prediction, encounter classification, and maritime traffic complexity analysis using large-scale AIS datasets [26]. These approaches demonstrate the potential of data-driven models to capture nonlinear relationships between environmental conditions, vessel characteristics, and traffic interactions. Nevertheless, relatively few studies have attempted to integrate environmental forcing, bathymetric constraints, and vessel-interaction dynamics into predictive models of pilot transfer behaviour. In particular, the spatial decision process governing pilot meeting-point selection remains largely unexplored in existing maritime research.

Insights from other transport domains demonstrate the potential benefits of dynamic rendezvous optimisation and trajectory-based decision support systems. In air-traffic management, continuous trajectory optimisation and mixed-integer nonlinear programming models have been developed to resolve aircraft conflicts under dynamic traffic conditions [33,34]. Decentralised airspace architectures have also explored adaptive conflict-resolution strategies within distributed traffic networks [35]. In land-transport and logistics systems, graph-based algorithms and spatio-temporal optimisation models have been used to determine optimal meeting points for vehicles operating in dynamic road networks [36–38]. Although these approaches illustrate the feasibility of dynamic meeting-point optimisation, they have not yet been systematically adapted to the maritime domain, where vessel manoeuvrability, hydrodynamic constraints, bathymetry, and regulated traffic-separation schemes introduce unique operational challenges.

The transition to maritime autonomous surface ships (MASS) adds a further driver: ports will need to redefine pilotage and handover arrangements [39], and the handover phase concentrates residual risk [40,41] (Section 6.3).

Taken together, the existing literature reveals several important gaps. First, environmental conditions, traffic interactions, bathymetric constraints, and vessel manoeuvring behaviour are typically analysed separately and presented in fragmented manner rather than within a unified operational framework. Second, most pilotage studies treat pilot boarding locations as static chart-based areas rather than dynamic

decision variables that may shift in response to operational conditions. Third, despite advances in AIS-based maritime analytics and machine-learning techniques, few studies have developed predictive frameworks capable of estimating pilot meeting distances using fused operational datasets. Fourth, although pilot-transfer accident inquiries commonly consider a small set of recurring risk categories (grounding, collision, allision, occupational injury and situational-awareness failure), the existing literature does not provide a structured mapping between data-driven decision variables and these categories.

3. Data and preprocessing

This study integrates a multi-source dataset covering one full operational year (September 2023–August 2024) for inbound and outbound pilotage at Busan Newport. Busan Newport was selected as the case-study port due to its high traffic density, complex approach geometry, and intensive pilotage operations associated with one of the largest container hubs in Northeast Asia. The port handles a large number of deep-draft container vessels operating under tidal constraints and narrow approach channels, creating frequent interactions between inbound, outbound, and coastal traffic streams. These operational characteristics make Busan Newport an appropriate test environment for analysing pilot meeting-point relocation behaviour under varying environmental and traffic conditions. Three complementary data streams including AIS trajectories, environmental and bathymetric observations, and traffic-interaction context, were time-synchronised at one-minute resolution and preprocessed into a unified analytical framework for modelling dynamic pilot meeting distances.

3.1. AIS trajectory data and reconstruction

AIS data constitute the backbone of the analysis and include vessel identity, gross tonnage (GT), timestamp, latitude/longitude, speed over ground (SOG), and course over ground (COG). Distances are reported in nautical miles (NM) where consistent with operational practice. After removal of physically implausible messages (e.g., duplicated timestamps, abnormal spatial jumps, negative or discontinuous speeds), vessel trajectories were reconstructed into continuous, time-ordered paths for all pilotage sessions. The AIS data supplied for the study cover 34.86–35.07°N and 128.67–128.88°E. This extent contains the approach channels and boarding zones but not the New Port berths or the pilot-boat base, which lie north of 35.07°N. Outbound ships and departing pilot boats therefore enter the data only shortly before an outbound transfer, which limits the decision lead time available for outbound events (Section 4.4).

Pilot boats were identified by MMSI. The fleet list originally supplied for the study also contained 44 harbour tugs; it was checked against vessel registers and AIS ship type, and only the 13 Busan pilot boats are used. Each AIS-detected transfer was then validated against the public pilotage log of the Busan Harbor Pilots' Association for the New Port pilot station (PS-3), which records the ship, the time and the direction of every pilotage act. A detected session is kept when the same ship (matched by call sign, otherwise by name) has a logged New Port act within ±2 h; the direction is taken from the log. Of 4599 detected sessions, 3243 are confirmed (precision 70.5%; 1664 inbound and 1579 outbound transfers on 961 ships), and these form the study dataset. All AIS timestamps are Korean Standard Time.

Fig. 1 shows the structure of the traffic around the approach: the density of AIS reports for one representative week (March 2024) on the official chart, with the tracks of the verified pilot boats in the same week. The high-density corridors of the approach channel, the coastal-transit lanes and the pilot-boat runs between the harbour and the boarding zones are clearly separated. All chart backgrounds in this paper are Korea Hydrographic and Oceanographic Agency (KHOA) electronic navigational charts KR5G3A44 and KR5G3E22 (harbour) and KR65886D, KR65887D, KR65896D and KR658D97 (berthing); areas

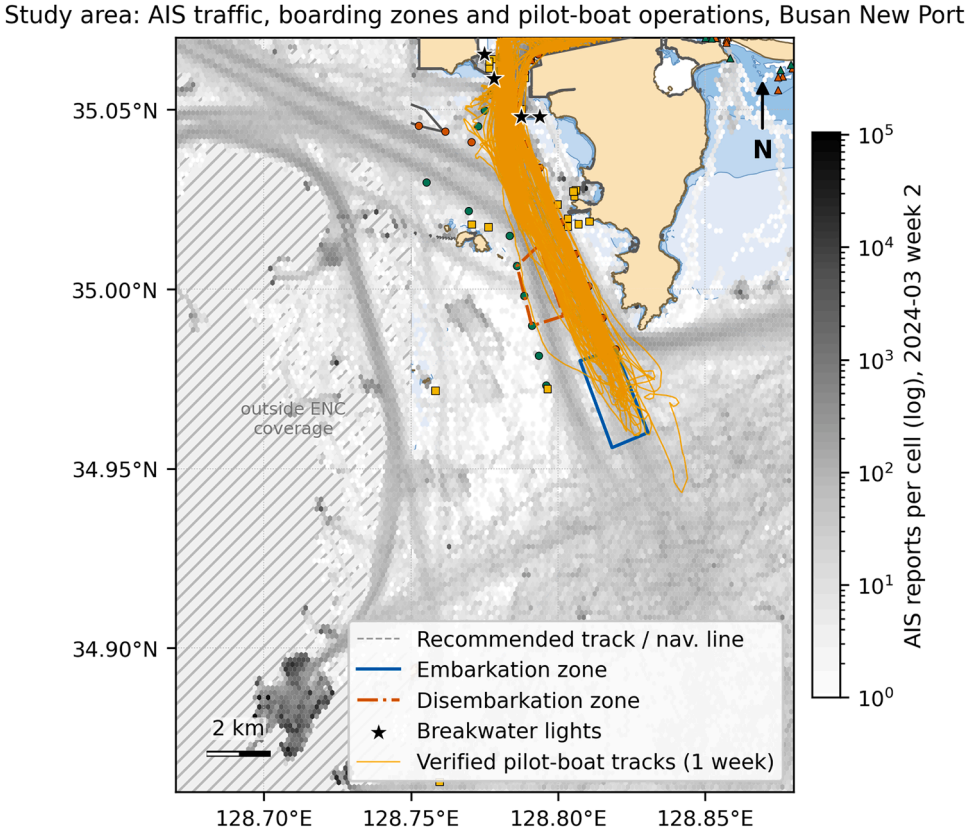

Fig. 1. AIS traffic density and verified pilot-boat tracks around Busan Newport, second week of March 2024.

outside their coverage, west of 128.75°E, are drawn from Natural Earth.

Spatial derivatives of AIS trajectories, including approach-distance profiles, COG variability, turning-rate indicators, and inbound/outbound speed-stability patterns, were extracted for each pilotage session. These derived variables characterise vessel motion behaviour in the vicinity of the designated pilot polygons and inform the model's ability to detect operationally unstable, high-workload approach geometries.

3.2. Environmental and bathymetric data

Meteorological and sea-state observations were obtained from coastal and offshore monitoring stations operated by the Korea Meteorological Administration and regional oceanographic buoy networks. The selected environmental variables include wind speed and direction, visibility class, significant wave height (H_s), wave period and direction, and hourly tidal elevation. These variables were selected because they represent the primary environmental factors influencing pilot transfer safety and vessel manoeuvring conditions, as highlighted in international pilotage guidance and previous maritime safety studies.

Hourly wind and wave observations from four Korea Meteorological Administration (KMA) networks, namely automatic weather stations (AWS), marine weather buoys, wave-height buoys and lighthouse weather units, were supplied for the study by GCS Solution Co., Ltd. The study region contains 22 listed stations (11 AWS, 6 wave-height buoys, 3 lighthouse units and 2 marine buoys; Fig. S2). Only 13 of them report wind and 8 report wave height, their records cover only January–August 2024, and none reports visibility. The full-year analysis therefore uses one consistent series, the Geoje marine weather buoy (KMA 22104), which is the only source spanning the whole study year; it lies a median 32.3 km (95th percentile 34.8 km) from the transfer positions. Short gaps were interpolated in time. The approximation error of this choice, and of nearest-station assignment and inverse-distance weighting where several stations exist, is quantified in Section 4.5. Visibility could not be included, because it is not recorded in the study area. Tide heights are predicted hourly values.

Water depth was taken from the Korean electronic navigational chart (ENC) cells covering the approach (depth areas and soundings). The General Bathymetric Chart of the Oceans (GEBCO) grid [42] was used only for visual context, because its ∼450 m resolution places real transfer positions in implausibly shallow water. Combined with tidal elevation data, vessel draught information, and local seabed slope, these datasets were integrated to estimate UKC for each vessel state vector. This integration enabled the identification of shallow-water constraints and bathymetric features that influence safe pilot-boat approach geometry and pilot transfer location selection.

3.3. Pilot-Interaction and traffic-context features

Traffic-interaction metrics were derived by synchronising AIS trajectories of all vessels operating within the port approaches using timestamp alignment and nearest-time interpolation to construct simultaneous traffic snapshots [23,43]. For each pilot--boat–subject-vessel pair, a traffic snapshot was generated at one-minute resolution, capturing the number of target vessels within a radius of 0.5–3 nautical miles around the subject vessel, encounter durations, minimum passing distances, and localised congestion intensity expressed as a three-level traffic density class.

Pilot transfer positions are extracted from the AIS interactions; the detection rules are given in Section 4.2. Detected sessions are then confirmed against the pilotage log as described in Section 3.1, which replaces the outlier filtering used in earlier versions.

The resulting cleaned inbound and outbound transfer points are shown in Fig. 2 together with the officially designated embarkation and disembarkation polygons. The spatial distribution reveals substantial dispersion around the formal pilot boarding zones, indicating that operational transfer locations frequently deviate from charted pilot-boarding areas in response to prevailing traffic and environmental conditions.

The polygons mark the regulated embarkation and disembarkation zones defined by the port authority. The dispersion of the transfer points around them is quantified in Section 5.2, where the zone centroid serves

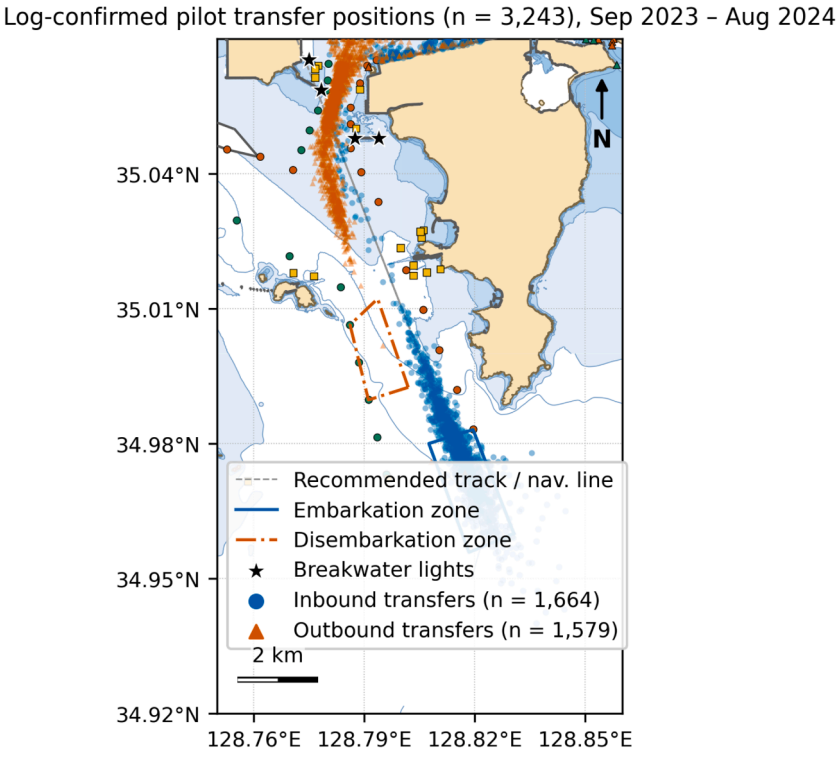

Fig. 2. Positions of the 3243 log-confirmed pilot transfers and the official boarding polygons.

as baseline B1.

3.4. Target variable definition

The target variable is the meeting distance: the distance from the transfer position, i.e. the point of minimum hull separation between the pilot boat and the subject vessel during the session, to the nearest New Port breakwater light. Inbound and outbound meeting distances are modelled separately, because they follow different geometries (Section 4.8).

3.5. Predictor variables and empirical distribution

All predictors are measured at the decision time t_d, the dispatch of the pilot boat (Section 4.4), using only observations made at or before t_d; the transfer time and anything observed later are used only as outcomes. Table 1 lists the predictors. Of the 3243 confirmed transfers, a dispatch could be identified for 2997 (1605 inbound and 1392 outbound), and these form the modelling dataset.

The inbound meeting distance (median 7.4 km from the breakwater) is concentrated along the approach channel, whereas outbound transfers (median 0.8 km) take place close to the breakwater with a much smaller spread (standard deviation 0.67 km against 2.74 km inbound). This difference supports separate inbound and outbound models (Section 4.8).

4. Methodology

This section describes the methodological framework developed to predict dynamic pilot meeting points and recommend relocation-adjusted rendezvous positions under varying operational conditions. The complete workflow, from raw AIS ingestion through feature fusion, model prediction, safety filtering, and dynamic waypoint recommendation, is summarised in Fig. 3.

Vessel naming convention.

Throughout the remainder of this paper the terminology follows the convention adopted in AIS-based encounter and pilotage analytics [21–23]: the *subject vessel* is the inbound or outbound ship engaged in pilot transfer; the *pilot boat* is the small craft delivering or collecting the pilot; and *target vessels* (or *surrounding traffic* when treated in aggregate) denote all other ships used to construct closest-point-of-approach (CPA), encounter, and traffic-density indicators around the subject vessel.

Table 1
Predictor variables at pilot-boat dispatch.

Category	Variable	Description
Ship state	Visibility	Whether the ship is already within the AIS extent at t_d
	Position	Distance (m) and bearing (sin, cos) from the breakwater
	Speed and course	SOG (kn); COG (sin, cos)
	Ship characteristics	Declared draught (m); AIS ship-type class
Pilot boat	Position	Distance (m) and bearing (sin, cos) from the breakwater
	Speed and course	SOG (kn); COG (sin, cos)
Relative	Range and bearing	Distance (m) and bearing (sin, cos) from the pilot boat to the ship
	Kinematic intercept	Distance from the breakwater and time (min) at which the pilot boat, at its speed (at least 12 kn), would meet the ship holding its course and speed
Traffic	Vessel counts	Third-party vessels within 0.5, 1 and 2 NM of the ship (service craft excluded)
	Nearest vessel	Distance to the nearest third-party vessel (m)
Metocean	Wind, waves	Wind speed (m/s) and H_s (m) of the hour containing t_d
Calendar	Hour, month	Hour of day (KST) and month of t_d

Generic terms such as “vessel manoeuvrability” or “vessel dimensions” retain their conventional meaning and refer to ship properties in general, independent of role.

4.1. Data preprocessing

All raw data sources, including AIS trajectories (1-min resolution), meteorological observations, sea-state measurements, tidal records, bathymetric grids, and vessel particulars, were harmonised into a unified spatiotemporal framework. The preprocessing stage ensured consistency across heterogeneous datasets before feature extraction and model development.

The preprocessing pipeline included several standardisation steps:

- alignment of all sources to Korean Standard Time (KST), the time base of the AIS, metocean and pilotage-log records
- removal of records with missing identity, time or position and of duplicated (MMSI, time) records
- normalisation of environmental variables, including wind speed, wave height, and tidal elevation, to consistent physical units
- coordinate validation and transformation into the WGS-84 geographic reference system (EPSG:4326)

Meteorological and sea-state values were matched to AIS times by hour. Depths were sampled from the ENC layers described in Section 3.2, ensuring consistent environmental and seabed attribution for all vessel positions. The resulting integrated dataset provided a temporally synchronised representation of vessel movements, environmental conditions, and bathymetric constraints for subsequent modelling and analysis.

4.2. AIS trajectory reconstruction and proximity event detection

Pilot-transfer sessions were detected from the synchronised trajectories of the 13 pilot boats and all other vessels, binned at one minute. Distances between positions are great-circle (haversine) distances [44]; hull separation uses the AIS antenna offsets and ship dimensions.

Session Detection. A pilot boat and a ship form a candidate pair in a minute bin when their hull separation is at most 500 m and the ship’s AIS dimensions or tonnage identify it as a vessel (length overall, LOA $\geq$ 60 m, beam >2 m or tonnage >500; in practice this removes only targets without dimensions). Consecutive bins of a pair form one session unless they are >15 min apart, and sessions shorter than 10 min are discarded. Three kinematic indicators are evaluated in every bin:

Course alignment

$$\left|COG_{ship} - COG_{pilot}\right| \leq \varepsilon_\theta, \qquad \varepsilon_\theta = 20^\circ$$

Speed similarity

$$V_{ship}, V_{pilot} \in [4, 12]\,\text{kn}, \qquad \left|V_{ship} - V_{pilot}\right| \leq \varepsilon_v, \qquad \varepsilon_v = 4\,\text{kn}$$

Boarding signature

$$V_{ship}, V_{pilot} \leq 2\,\text{kn}, \qquad \max\left(V_{ship}, V_{pilot}\right) > 0.3\,\text{kn}$$

A session is retained when course alignment or speed similarity holds in at least half of its bins, or the boarding signature in at least 30%. The transfer position and time are those of the minimum hull separation in the session, and the direction is assigned from the ship’s movement relative to the port (later replaced by the logged direction). All retained sessions are then checked against the pilotage log (Section 3.1); only the 3243 confirmed transfers are used. The log also measures the detector’s coverage: the confirmed sessions account for 28.4% of logged inbound and 26.6% of logged outbound New Port acts, and the share is almost unchanged (28.8% and 27.2%) when restricted to ships visible in the AIS. Missed transfers therefore come from the detection rules rather than the AIS extent, and the confirmed set is a validated sample rather than a

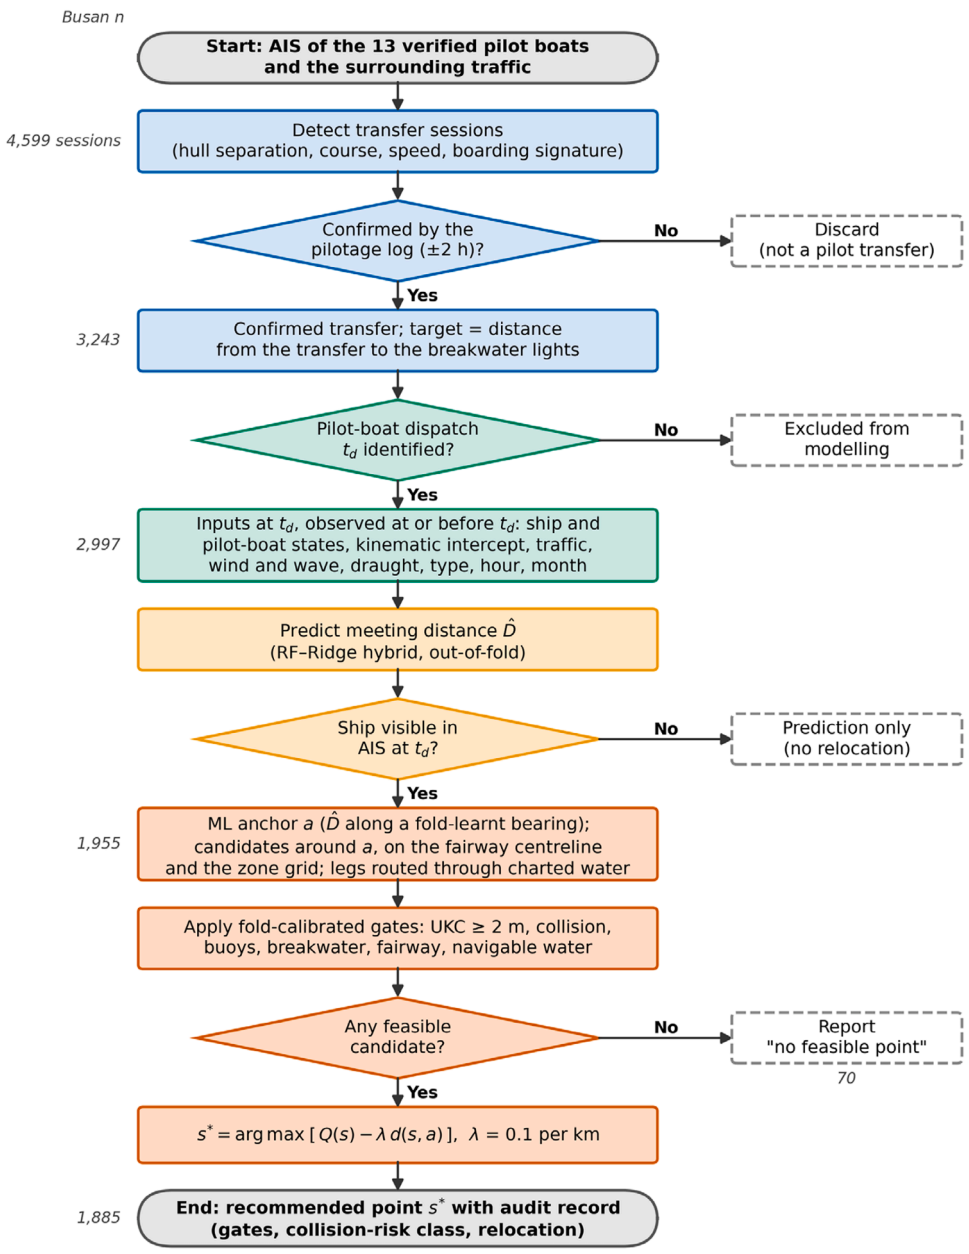

Fig. 3. Workflow of the framework as an algorithm flowchart (Busan counts on the left).

census of all transfers.

4.3. Computation of meeting distance points (Target variable)

The predicted variable is the meeting distance,

$$D_{meet} = \min_{k} d((\phi_m, \lambda_m,\ \phi_k, \lambda_k)),$$

where (ϕ_m, λ_m) is the transfer position and (ϕ_k, λ_k) are the four New Port breakwater lights. The model predicts this scalar distance; the position along which it is laid out is determined separately in the optimisation (Section 4.10). Inbound and outbound distances are modelled separately.

4.4. Decision time

A meeting point has to be chosen before the transfer, at a moment that is observable when the decision is made. We use the dispatch of the pilot boat, the moment at which the pilot station sends the boat to a known ship. It is identified in the AIS as follows. After the same pilot boat has ended its previous transfer session (or, at most, 90 min before the transfer), the dispatch time t_d is the boat's first fix with SOG above 5 kn whose course is within 60° of the bearing to the ship or, if the ship is not yet visible, points seaward away from the breakwater. The dispatch is thus a proxy identified retrospectively: the ship assignment is taken from the detected session, and the search window is bounded by the observed transfer (at most 90 min before it). The observed transfer time T is never an input, and no observation made after t_d is used (each input timestamp is checked to be at or before t_d).

A dispatch is found for 1605 of 1664 inbound and 1392 of 1579 outbound transfers. The median time from dispatch to transfer is 22.7 min inbound and 24.7 min outbound. At dispatch the ship is already visible in the AIS for 73.7% of inbound and 56.1% of outbound transfers; for the others it is still beyond the AIS extent (Section 3.1), which is itself known at t_d and enters the model as a predictor.

The definition of the decision point matters, and we examined it

carefully. Defining it relative to the observed transfer (a fixed time before T) is not possible online and inflates accuracy, because the ship's state a known time before the meeting nearly determines where it will be. Taking the first movement after the boat's *last* wait before the session, rather than its first movement towards the ship, also uses hindsight, because only later observations reveal which wait is the last; it raised inbound R^2 from 0.495 to 0.645. Conversely, the ship's first AIS fix within 13 km of the breakwater is observable but carries no information about when the boat will be sent, and gives no predictive skill ($R^2 \approx 0$). Finally, the 90 min bound itself uses the transfer time. Without it, starting the search only at the end of the boat's previous detected session, the dispatch changes for 39% of the transfers, the median time from dispatch to transfer rises to 64 min and inbound R^2 falls to 0.30 (temporal hold-out 0.32; outbound 0.04; 1397 inbound and 1269 outbound transfers with a dispatch under this rule, so the comparison is not on identical events). Because the detector finds only part of the pilot boats' jobs, the previous detected session often lies hours back, and the search then picks up movements for other jobs. The results below therefore describe prediction at a dispatch proxy within 90 min of the transfer, evaluated retrospectively; with a logged dispatch time the bound would not be needed.

4.5. Environmental field construction

For every event, wind speed and significant wave height are taken from the hourly Geoje buoy series at the hour of the decision time (Section 3.2). Because this buoy lies about 32 km offshore, we quantified the approximation error with the multi-station data available for January–August 2024. In leave-one-station-out cross-validation, inverse-distance weighting (IDW, power 2),

$$\widehat{z}(\boldsymbol{x}_0, t) = \frac{\sum_s w_s z_s(t)}{\sum_s w_s}, \qquad w_s = d(\boldsymbol{x}_0, \boldsymbol{x}_s)^{-2}, \tag{1}$$

was more accurate than nearest-station assignment, $z_{s^*}(t)$ with $s^* = \mathrm{argmin}_s d(\boldsymbol{x}_0, \boldsymbol{x}_s)$: root mean square error (RMSE) 2.32 vs 2.75 m/s for wind (13 stations) and 0.36 vs 0.41 m for H_s (8 stations), with IDW better at 69% and 75% of the stations. Re-fitting the meeting-distance model on the 2024 events with nearest-station or IDW inputs instead of the buoy series changed R^2 by 0.000 inbound ($n = 1016$) and by -0.008 to -0.009 outbound ($n = 522$), with mean absolute error (MAE) changes below 5 m. The model is therefore insensitive to the assignment scheme, because its skill comes mainly from the ship's state at decision time. The single-buoy series is retained for full-year consistency, and IDW is recommended where full-year multi-station data exist.

4.6. Under-keel clearance (UKC) computation

UKC was computed for each vessel state vector using bathymetric and tidal information.

$$\mathrm{UKC} = D_{\mathrm{chart}} + T_{\mathrm{tide}} - d_{\mathrm{draft}} \tag{2}$$

where

- D_{chart} represents charted depth (positive downwards) from the ENC depth areas and soundings
- T_{tide} denotes the predicted tide height at the hour of the evaluated time
- d_{draft} corresponds to the vessel draught reported in AIS or ship particulars.

Dynamic squat. Eq. (2) gives static UKC. A ship under way also sinks and trims by squat, which grows with the square of speed through the water and with block coefficient C_B, and is amplified in shallow and confined water by the blockage factor. Using Barrass's estimate, $S_{\max} \approx C_B V^2/100$ in open water and about twice that in confined water (V in kn, S in m), with $C_B = 0.65$ for container and cargo ships and 0.80 for tankers, the ships in the dataset (median speed 8.5 kn at the transfer, interquartile range 7.1–10.1 kn) have a median squat of 0.47 m in open and 0.94 m in confined water (95th percentile 1.87 m). Subtracting confined-water squat lowers the share of real transfers with UKC ≥ 2 m from 99.6% to 98.3% (of the 3230 transfers with a recorded speed), and the share of the framework's recommendations from 100% to 98.9%. Squat is therefore reported as a sensitivity rather than built into the constraint, because ship-specific coefficients and speed through the water are not available.

Role of UKC. UKC is used as a hard safety constraint in the optimisation (Section 4.10), evaluated at every candidate point. It is not a predictor of the meeting distance. UKC at the observed meeting point can be computed only after the meeting position is known, so as a predictor it would leak the target; a decision-time alternative is tested in Section 5.2. The constraint level of 2 m is deliberately more conservative than the one local numerical rule we found, in the Busan New Port passage plan published by the Busan Harbor Pilots' Association (a minimum UKC of 0.5 m at berth and 7% of the ship's deepest draught in the fairway, i.e. about 1.0 m for a 15 m draught) [45], and includes a margin for squat and depth uncertainty. International pilotage guidance gives no numerical limits for wind, waves or visibility, so these are recorded but not used as constraints.

4.7. Dynamic traffic proximity modelling

Traffic complexity was quantified from synchronised AIS observations of all target vessels in the study area, excluding service craft (tugs, pilot boats and port tenders by AIS ship type), and counting each vessel once at its latest position within the evaluation window. For each pilot-transfer timestamp, a dynamic traffic snapshot was constructed describing the spatial configuration and encounter characteristics of the surrounding traffic.

Local target-vessel density was computed as

$$\rho(t) = \frac{N_r(t)}{\pi r^2} \tag{3}$$

where $N_r(t)$ denotes the number of target vessels located within a circular neighbourhood of radius r centred at the subject vessel. The target-vessel count is defined as

$$N_r(t) = \sum_i 1[d_i(t) \leq r] \tag{4}$$

where $d_i(t)$ represents the distance between the subject vessel and target vessel i at time t, and $1[\cdot]$ is an indicator function equal to 1 when the condition is satisfied and 0 otherwise.

In addition to density, several encounter-level metrics were computed, including the number of proximate target vessels within 1–2 nautical miles, encounter duration, traffic directionality (inbound/outbound classification), and closest-point-of-approach (CPA) proximity measures.

Conventional AIS-based CPA is computed between the geometric reference points of two vessels and therefore treats each ship as a point object. This assumption can lead to optimistic estimates of separation distance, particularly in confined waterways or during pilot boarding operations where vessel dimensions are significant relative to encounter distance. To address this limitation, a dimension-corrected proximity metric, denoted $\mathrm{CPA}_{\mathrm{SD}}$, was introduced.

We distinguish two quantities. The signed size-corrected separation

$$\Delta_{\mathrm{SD}} = D_{\mathrm{CPA}} - 1/2(L_O + L_T) \tag{5}$$

can be negative, and the net clearance

$$\mathrm{CPA}_{\mathrm{SD}} = \max(0,\ \Delta_{\mathrm{SD}}) \tag{6}$$

cannot. D_{CPA} is the centre-to-centre distance at the time of closest approach (TCPA, set to zero when the closest approach is already past) between the subject vessel and the nearest target vessel, and L_O and L_T are their lengths. Subtracting the half-lengths approximates the longitudinal extent of both hulls. $\Delta_{SD} < 0$ means that the ships came within their combined length envelope; it does not by itself show that the hulls overlapped, because the length-only correction overstates the hull extent whenever the closest approach is not along both ships' long axes, which in our data applies to crossing and head-on as well as side-by-side encounters (Section 5.4). It is therefore used as a conservative screening approximation, and we also compute (i) a beam-only correction, (ii) a heading-weighted projection $1/2(L|\cos\theta| + B|\sin\theta|)$ per ship, with θ the angle between the ship's heading and the line of sight, and (iii) the minimum distance between oriented hull rectangles ($L \times B$, positioned with the AIS antenna offsets and oriented by heading, or by course when heading is unavailable). Encounters are classified by relative heading as head-on ($|\Delta\psi| \geq 165^\circ$), overtaking ($|\Delta\psi| \leq 67.5^\circ$) or crossing, and as stationary when either ship moves below 1 kn.

In the optimisation (Section 4.10), the collision-risk class of a candidate point uses the domain clearance: the centre CPA minus the two ship-domain radii, each the hull half-diagonal plus $\max(0.25L, 1.5B)$. A clearance below 100 m is high risk, below 300 m medium and otherwise low. The same formulation could be extended to ship-to-fixed-infrastructure proximity (breakwater, navigation beacons, channel edge) by replacing L_T with the projected dimension of the obstacle, which would support allision-risk quantification.

4.8. Hybrid machine-learning model

The predictive architecture combines Random Forest regression with a Ridge residual-correction module. This hybrid design enables the model to capture nonlinear operational relationships while maintaining stable behaviour in data-sparse regimes.

Random Forest was selected as the primary learning algorithm because it performs well on heterogeneous maritime datasets containing nonlinear interactions between environmental variables, vessel characteristics, and traffic conditions. Tree-based ensemble models are also robust to multicollinearity and require minimal distributional assumptions, making them suitable for AIS-derived operational data.

Its performance is compared with six baselines under identical partitions and preprocessing in Section 5.2.

Hyperparameters are fixed and not tuned on the data: 400 trees, unrestricted depth, a minimum of 2 samples per leaf, and Ridge $\alpha = 1$ for the residual stage. Predictors are median-imputed and categorical variables one-hot encoded, with all preprocessing fitted on the training data only.

A Ridge residual-correction layer was added with the aim of smoother behaviour in sparse operational regimes; we did not test this separately, and in the results it changes accuracy little.

Stage 1: Random Forest Regression. Let the feature vector describing a pilot-transfer event be

$$\boldsymbol{x}_t = [x_{1,t}, x_{2,t}, \ldots, x_{p,t}]$$

The Random Forest model predicts the meeting distance as

$$\widehat{D}_{RF,t} = f_{RF}(\boldsymbol{x}_t) \tag{7}$$

where $f_{RF}(\cdot)$ denotes the fitted Random Forest estimator.

Stage 2: Ridge Residual Correction. Residuals from the Random Forest model are defined as

$$e_t = D_{meet,t} - \widehat{D}_{RF,t} \tag{8}$$

A Ridge regression model is then trained to approximate these residuals:

$$\widehat{e}_t = g_{Ridge}(\boldsymbol{x}_t) \tag{9}$$

The final hybrid prediction becomes

$$\widehat{D}_t = \widehat{D}_{RF,t} + \widehat{e}_t \tag{10}$$

Inbound and Outbound Model Separation. Separate models were trained for inbound and outbound operations because the two scenarios exhibit fundamentally different dynamics:

- inbound subject vessels follow a constrained approach corridor.
- outbound subject vessels leave from the berths over a broad departure sector.
- environmental exposure differs between the two phases.
- traffic interaction patterns vary significantly.

A hybrid Random Forest–Ridge regressor is fitted for each direction. The submitted version of this paper modelled outbound transfers with a route-stratified ensemble, because the outbound meeting points then formed two bearing clusters (near 113° and 238° from the breakwater). On the verified pilot-boat data this structure is absent: outbound meeting bearings are unimodal (Section 5.2), and the earlier south-east cluster was formed by harbour-tug interactions that the corrected fleet list removes. We therefore use the single hybrid for outbound transfers as well, and report the route-stratified ensemble, with its route classifier and an oracle variant, only as an alternative in Section 5.2.

In the revised results the residual stage changes accuracy little (the hybrid and the random forest perform almost identically, Section 5.2); we retain it as a regularised correction, without claiming better behaviour in sparse regimes.

4.9. Model validation

Three partition schemes are used. (i) Random 5-fold cross-validation over events, for comparison with the submitted version. (ii) Ship-grouped 5-fold cross-validation (GroupKFold by subject-vessel MMSI), so that no ship appears in both training and test data; this is the primary scheme, because many ships call repeatedly. (iii) A forward temporal hold-out: training on transfers before 1 July 2024 and testing on July–August 2024 (190 inbound and 132 outbound test events). In every scheme imputation, encoding, any clustering and any classifier are fitted on the training part only, and hyperparameters are fixed. Metrics are computed on the pooled out-of-fold predictions, with the standard deviation across folds; 95% intervals are obtained by a bootstrap that resamples whole ships (1000 resamples), because repeated calls of one ship are not independent, and are conditional on the fitted folds.

Model accuracy was assessed with root mean square error (RMSE), mean absolute error (MAE), the coefficient of determination (R^2) [46–48]; no calibrated prediction interval is evaluated. Empirical error ranges are taken from the quantiles of the pooled out-of-fold residuals; they describe the cross-validated error distribution and are not a separately calibrated prediction interval.

Six baselines are evaluated under the same partitions, features and preprocessing: (B1) the official fixed boarding zone, predicting the distance of the embarkation or disembarkation polygon centroid; (B2) the historical mean distance by direction; (B3) linear regression; (B4) random forest without residual correction; (B5) gradient boosting, both scikit-learn HistGradientBoosting and extreme gradient boosting (XGBoost; 400 trees, learning rate 0.05, depth 6, subsampling 0.8, untuned); and (B6) the proposed hybrid.

4.10. Safety-constrained meeting-point optimisation

The model predicts a distance, not a position, and it ignores hard safety limits. The optimisation turns the prediction into a recommended position at dispatch and relocates it only where safety requires. It uses

only information available at t_d: the traffic picture is the latest AIS report of each vessel within the 5 min before t_d, the positions of the ship and the pilot boat are their latest fixes at or before t_d, and no part of either craft's later track is used. The optimisation is run for the transfers in which the ship is already visible at dispatch, because its position is needed.

Anchor. For each event the predicted distance $\widehat{D}$ is the out-of-fold prediction, i.e. from a model that did not see that ship. The anchor a is the point at distance $\widehat{D}$ from the breakwater lights along an anchor bearing learnt from the training fold only: the circular mean of the meeting bearings of the 25 training transfers whose ships had the most similar bearing from the breakwater at dispatch.

Candidates and routes. The candidate set $\mathscr{S}$ contains the anchor; a ring of points 150, 300 and 500 m from it at 45° steps; a 300 m grid inside the official boarding polygon (at most 36 points); fairway-centreline points between paired lateral buoys near the anchor; points along and beside the ship's line to the anchor; the median position of earlier transfers in the same direction; and the midpoint of the pilot boat and the ship at t_d. The legs of the pilot boat and of the ship to each candidate are routed through charted water: the shortest path on a 100 m grid (200 m for long legs) from which all ENC land cells are removed. These legs, not straight lines, are tested against the route gates.

Constraints. A candidate s is feasible when

$$\mathrm{UKC}(s) \geq 2\,\mathrm{m} \tag{11}$$

$$\delta_{\mathrm{dom}}(s) \geq \delta_{\min} \tag{12}$$

$$d_{\mathrm{buoy}}(s) \geq 120\,\mathrm{m}, \quad d_{\mathrm{breakwater}}(s) \geq 250\,\mathrm{m} \tag{13}$$

$$\text{fairway access}(s), \text{ both legs clear of buoys and within navigable water} \tag{14}$$

and, for inbound transfers, the point lies ahead of the ship along its course. UKC uses Eq. (2) with the ship's draught and the predicted tide at t_d, and δ_{dom} is the domain clearance to the nearest target vessel at t_d (Section 4.7). UKC ≥ 2 m is a fixed rule. The navigational and traffic thresholds are empirical accommodation thresholds rather than proven safety limits: each is relaxed from its default only as far as needed for 95% of the completed transfers to satisfy it, and it is calibrated separately for each cross-validation fold, excluding every transfer (in either direction, with or without a dispatch) of the ships held out in that fold, so that no evaluated ship informs its own gates. The folds are grouped by ship, not by time, so a fold's gates may draw on transfers from later dates; for the training transfers the calibration uses their realised tracks, which is legitimate for training data but means that the gates describe completed practice rather than a forward-validated standard. Across folds the calibrated values are $\delta_{\min}$ from −272 to − 263 m (negative, because the ship-domain radii are generous and real transfers routinely take place inside them), fairway-edge distance 1318–1444 m and route clearance from buoys 18.66–20.15 m; the buoy, breakwater and water-coverage gates keep their defaults. Passing these gates means that a point is no tighter than nearly all transfers that pilots actually carried out; it is not a proof of safety. The collision-risk class itself is not a hard gate: it enters the safety score, so a feasible recommendation can still lie in the medium or high class. Sea state and visibility are not used as point constraints, because the single wind and wave series does not vary across candidates and visibility is not recorded.

Objective. Each candidate receives a safety score

$$Q(s) = \sum_i w_i\, q_i(s), \tag{15}$$

where each component $q_i \in [0,1]$ and the weights are: UKC margin (UKC/8 m, capped; 0.24), collision class (0.15), traffic count within 0.5 NM (0.14), fairway membership (0.09) and depth inside it (0.10), progress ahead of the ship (0.08), clearance from buoys (0.06) and from the breakwater (0.05), navigational position (0.04), route length (0.04), route hazard (0.05), route clearance (0.03), route water coverage (0.02) and weather (0.05). The weights sum to 1.14, so Q ranges from 0 to 1.14; they are expert-set, not fitted. The recommended point is

$$s^* = \arg\max_{s \in \mathscr{S}_{\mathrm{feasible}}} \left[Q(s) - \lambda\, \frac{d(s,a)}{1\,\mathrm{km}} \right], \qquad \lambda = 0.1, \tag{16}$$

so that a move of 1 km away from the predicted point must raise the safety score by at least 0.1. The objective trades closeness to the predicted point against the score; it does not guarantee the shortest feasible relocation. If no candidate satisfies all constraints, the framework returns *no feasible recommendation* for that event; the best-scoring infeasible point is kept only as a diagnostic and is not counted as a recommendation. The static evaluation at t_d does not model how the traffic will have moved when the two craft reach the point, which we state as a limitation.

This formulation subordinates the ML prediction to the safety constraints: the predicted point is kept when it is feasible and not materially less safe than its neighbours, and moved when it is not.

To provide model interpretability beyond global feature-importance measures, SHapley Additive exPlanations (SHAP) values were computed for the Random Forest models. SHAP provides a theoretically grounded method for decomposing individual predictions into additive feature contributions based on cooperative game theory [49]. For a given prediction $\widehat{D}(x)$, SHAP attributes to each feature x_j a contribution ϕ_j relative to the expected model prediction:

$$\widehat{D}(x) - E[\widehat{D}] = \sum_j \phi_j(x)$$

where $E[\widehat{D}]$ represents the mean model prediction across the training dataset. This formulation ensures local additivity and allows each prediction to be interpreted as the sum of feature-specific effects. SHAP values therefore enable both global and local interpretation of model behaviour and are used in this study to analyse how environmental conditions, vessel characteristics, and traffic interactions influence predicted pilot meeting distances.

5. Results

5.1. Optimisation output summary

Table 2 traces the confirmed transfers through the framework. A pilot-boat dispatch was identified for 1605 of the 1664 inbound and

Table 2
Data flow and optimisation outcome by direction.

Stage	Inbound	Outbound	Total
AIS-detected transfer sessions			4599
Confirmed by the pilotage log	1664	1579	3243
Pilot-boat dispatch identified (modelling dataset)	1605	1392	2997
Ship visible in AIS at dispatch	1183	781	1964
Optimisation attempted	1177	778	1955
No feasible recommendation	64	6	70
Feasible recommendation	1113	772	1885 (96.4%)
of which UKC ≥ 2 m	1113	772	1885 (100%)
of which low collision-risk class	969	598	1567 (83.1%)
Median time from dispatch to transfer (min)[a]	17.9	4.8	
Median distance, anchor to observed point (m)[a]	1129	1618	
Median relocation from the anchor (m)[a]	300	300	
Median distance, recommendation to observed point (m)[a]	1022	1323	1133

[a] Over the 1885 feasible recommendations.

1392 of the 1579 outbound transfers; for the others the pilot boat shows no dispatch movement towards the ship in the AIS (most often because it was already close to the ship or appears only at the meeting). At dispatch the ship was already visible for 1183 inbound and 781 outbound transfers, and the optimisation was run for 1955 of these (9 lacked a usable traffic picture). Among the confirmed transfers wind is missing for 20, wave height for 30, tide for 8 and depth for 2 events; these are imputed within the training folds.

Constraints at the predicted point. Of the 1931 ML anchors with a complete gate evaluation, 1019 (52.8%) satisfy every gate. The others fail on the buoyed route corridor (396), the route leaving navigable water (387), UKC below 2 m (298), fairway access (137), buoy clearance (44), collision risk (32) or the ship being already past the point (13); one anchor can fail several gates. UKC is active mainly outbound, where 211 of 752 anchors fall in shallow water close to the breakwater, against 87 of 1172 inbound.

Recommendations. Fig. 4 summarises the 1885 feasible recommendations. The recommended point is the anchor itself in 234 cases, a point on the 150–500 m ring around it in 1011, a fairway-centreline point between paired lateral buoys in 421, a point of the official zone grid in 125, a point in the ship's forward corridor in 44, a historical transfer hotspot in 35 and another candidate in 15. The median relocation from the anchor is 300 m in both directions. All recommendations satisfy UKC ≥ 2 m (98.9% after subtracting confined-water squat at the ship's speed at the meeting), and 1567 (83.1%) are in the low collision-risk class (medium 147, high 171). For 70 events (3.6%) no candidate satisfies all constraints, mostly because every route to the predicted area crosses the buoyed approach corridor (65); the framework then reports that no feasible point exists.

Agreement with practice. The recommended point lies a median 1022 m (inbound) and 1323 m (outbound) from where the transfer actually took place, and within 500 m of it in 17.9% and 18.4% of cases. These distances are not errors in the usual sense, because the recommendation is a feasible point near the predicted one under the implemented constraints and pilots may legitimately choose differently; they are of the same order as the prediction error itself (Section 5.2), and the relocation does not move the recommendation further from practice than the anchor (anchor to observed point: 1129 m and 1618 m). The relocation penalty λ sets the balance between closeness and the safety score (Fig. S3): without it ($\lambda = 0$) the optimiser seeks the highest-scoring point anywhere in the candidate set and the median distance to the observed point grows to 1391 m, whereas for $\lambda \geq 0.2$ the anchor is kept in most cases and the mean safety score falls. $\lambda = 0.1$ keeps most of the score gain at a median relocation of 300 m (median distance to the observed point 1133 m; anchor kept in 12.4% of cases).

5.2. Model performance evaluation

Leakage control. All results below use the pilot-boat dispatch proxy as the decision point (Section 4.4) and only observations made at or before it; the transfer time is never an input, but it bounds the dispatch search, so the evaluation is retrospective.

Accuracy. Table 3 reports the proposed hybrid model under the three partition schemes, and Fig. 5 shows its out-of-fold predictions under ship-grouped cross-validation. For inbound transfers ($n = 1605$ on 717

Table 3
Performance of the hybrid model by partition scheme (R^2 with fold standard deviation and, for the ship-grouped scheme, 95% interval).

Direction	Scheme	Train / test	R^2	MAE (km)	RMSE (km)
Inbound	Random 5-fold	1284 / 321	0.506 (0.028)	1.266	1.926
	Ship-grouped 5-fold	1284 / 321	0.495 (0.052) [0.442, 0.546]	1.278	1.948
	Temporal hold-out	1415 / 190	0.537	1.232	1.773
Outbound	Random 5-fold	1113 / 278	0.144 (0.052)	0.431	0.623
	Ship-grouped 5-fold	1113 / 278	0.159 (0.038) [0.082, 0.222]	0.429	0.618
	Temporal hold-out	1260 / 132	0.027	0.450	0.616

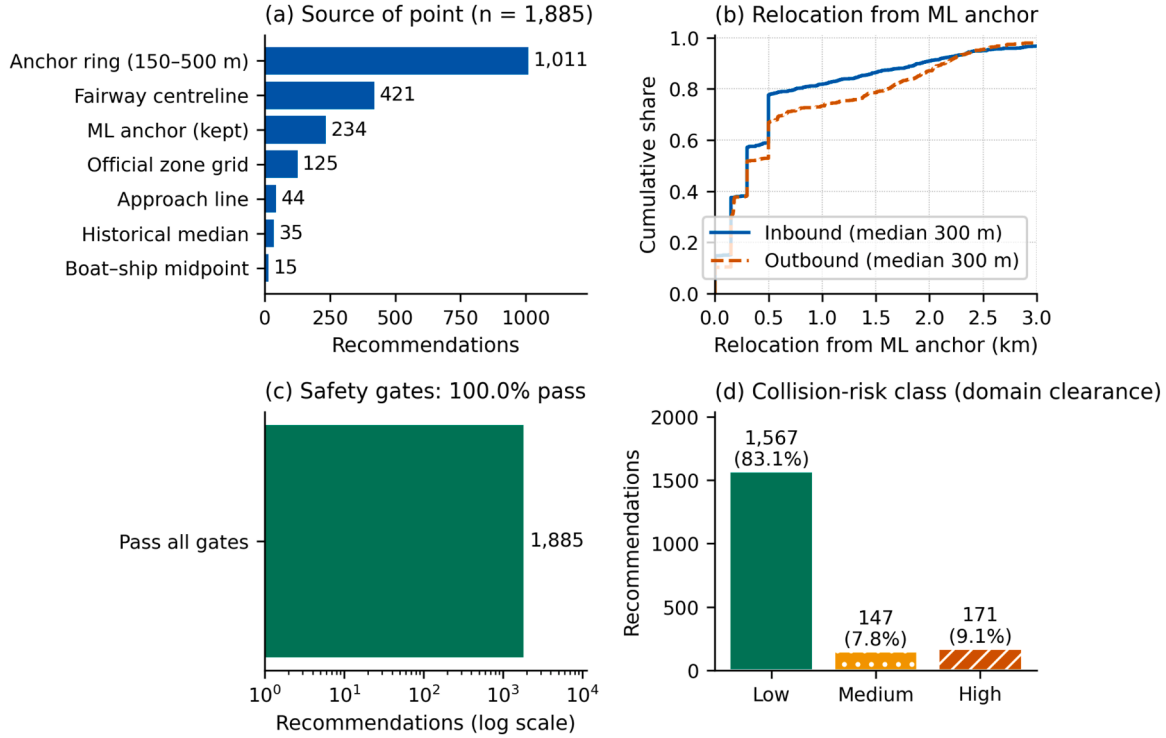

Fig. 4. Optimisation outputs: (a) Source of the recommended point; (b) Relocation from the ML anchor; (c) Safety-gate outcome; (d) Collision-risk class.

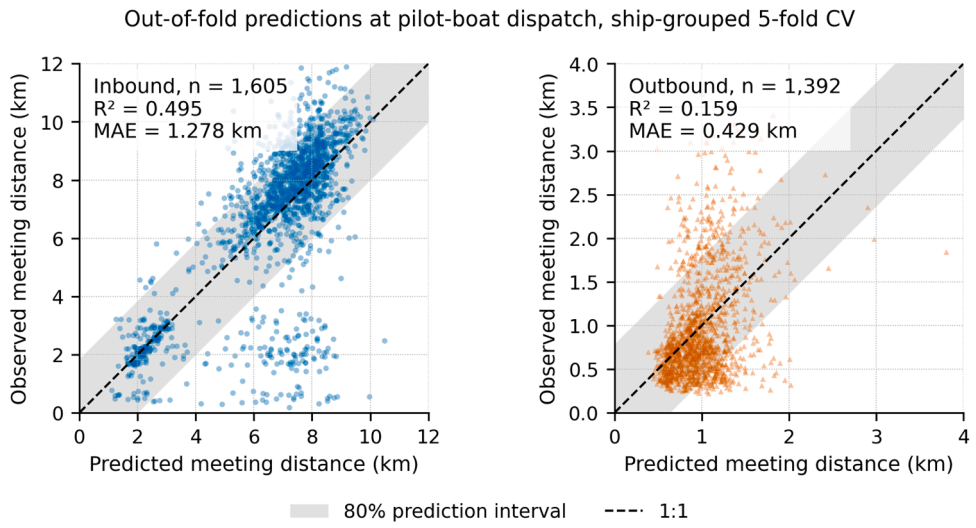

Fig. 5. Predicted versus observed meeting distance (ship-grouped out-of-fold): (A) Inbound; (B) Outbound.

ships), $R^2 = 0.495$ (95% ship-clustered confidence interval, CI, 0.442–0.546) with MAE = 1.278 km under ship-grouped cross-validation, and $R^2 = 0.537$ (MAE = 1.232 km) in the forward temporal hold-out, so the skill carries over to later months. When the ship is already visible at dispatch, R^2 is 0.56.

For outbound transfers ($n = 1392$ on 701 ships) the model reaches $R^2 = 0.159$ (0.082–0.222) with MAE = 0.429 km under ship-grouped validation and has no skill forward in time ($R^2 = 0.03$ in the temporal hold-out). Outbound pilot boats are often sent before the departing ship appears in the AIS (44% of cases), and when it is visible they are sent only minutes before the disembarkation; nothing observable at that point anticipates where it will take place. The absolute outbound error stays small because outbound transfers vary little (standard deviation 0.67 km). We therefore treat outbound prediction as not supported by the AIS data and use the outbound anchor only as a starting point for the safety check.

Table 4 shows how inbound skill falls with the time between dispatch and transfer: the earlier the boat is sent, the more the ship's subsequent speed and the pilots' choices remain open. Because this time becomes known only at the transfer, the table describes where the errors arise and is not an operating mode of the model.

Baselines. Table 5 compares the proposed model with the baselines under the same partitions. The official fixed boarding zone is by far the worst predictor ($R^2 = -0.70$ inbound and −42 outbound). This shows that a single fixed position does not describe where transfers take place; it does not by itself show that the fixed zones are unsafe. The historical mean has no skill. Among the learned models, the hybrid and the random forest are best and practically equal (inbound 0.495 and 0.494, outbound 0.159 and 0.152), with linear regression, gradient boosting and XGBoost slightly lower; inbound, the temporal hold-out gives the same picture, and outbound no model has meaningful skill in the temporal hold-out. We retain the hybrid for its stable performance across schemes and its exact SHAP decomposition.

Outbound route structure. The submitted version modelled outbound transfers with a route-stratified ensemble built on two bearing clusters near 113° and 238° from the breakwater. On the verified data the outbound meeting bearings form a single broad south-westerly sector (about 195–270°); the earlier south-east cluster was formed by harbour-tug interactions that the corrected fleet list removes. Splitting this sector into $k = 2$ sub-directions (silhouette 0.61) would raise outbound R^2 to 0.40 if the sub-direction were known in advance (oracle), but it cannot be predicted at dispatch (classifier-routed $R^2 \approx 0$), and $k = 3$ or 4 does not help (classifier-routed R^2 0.03 and − 0.02). For the same reason the mean distance per sub-direction ($R^2 = 0.32$) is not a usable baseline: it uses the sub-direction of the observed meeting point, which is known only after the transfer. The single hybrid is therefore used.

Error ranges. The 80% range of the out-of-fold errors (shaded in Fig. 5) is −1.85 to + 2.00 km inbound and −0.79 to + 0.63 km outbound; 35.8% of inbound and 68.7% of outbound predictions lie within 0.5 km of the observed distance, and 58.8% and 92.2% within 1 km. These ranges describe the cross-validated errors and are not a calibrated prediction interval.

UKC as a predictor. Adding decision-time UKC features (UKC at the ship's position from its draught, the predicted tide and ENC depth, and the distances from the breakwater beyond which UKC stays ≥ 2 m and ≥ 4 m along the approach line) changes R^2 from 0.495 to 0.499 inbound and from 0.159 to 0.158 outbound under ship-grouped validation. Depth rarely limits where pilots board at Busan Newport: median UKC at real transfers is 10.9 m and 99.5% have UKC ≥ 2 m. UKC is therefore

Table 4
Inbound accuracy by time from dispatch to transfer.

Dispatch before transfer	n	R^2	MAE (km)
0–5 min	130	0.90	0.51
5–15 min	332	0.81	0.80
15–30 min	515	0.68	1.00
30–90 min	628	0.17	1.91

Table 5
Baseline comparison (R^2; MAE in km in brackets).

Model	Inbound		Outbound	
	Ship-grouped	Temporal	Ship-grouped	Temporal
B1 Official fixed zone	−0.699 (2.564)	−0.962 (2.713)	−41.8 (4.365)	−48.4 (4.343)
B2 Historical mean	−0.000 (2.201)	−0.013 (2.100)	−0.001 (0.485)	−0.001 (0.468)
B3 Linear regression	0.475 (1.364)	0.496 (1.341)	0.107 (0.447)	0.048 (0.453)
B4 Random forest	0.494 (1.284)	0.539 (1.223)	0.152 (0.426)	0.055 (0.444)
B5a Gradient boosting (HistGB)	0.462 (1.357)	0.463 (1.359)	0.076 (0.451)	−0.120 (0.476)
B5b XGBoost	0.470 (1.323)	0.507 (1.309)	0.099 (0.438)	−0.168 (0.470)
B6 Proposed RF–Ridge hybrid	0.495 (1.278)	0.537 (1.232)	0.159 (0.429)	0.027 (0.450)

used as a hard constraint in the optimisation, where it is active (Section 5.1), not as a predictor.

Computational cost. On an Apple M4 Pro laptop (24 GB RAM; Python 3.12, scikit-learn 1.8), training takes under 0.5 s per direction and inference about 28 ms per event. The safety-constrained optimisation, including the water-routed legs, takes about 4 s per event. Run-time cost is therefore small compared with the time between dispatch and transfer; only data fusion and retraining are batch processes.

5.3. Feature importance analysis

SHAP values were computed for the hybrid models fitted on all modelling events, combining tree SHAP for the Random Forest stage with the exact linear SHAP of the Ridge residual stage, so that the attributions add up to the hybrid prediction (Fig. 6).

For inbound transfers the kinematic intercept dominates (mean $|\text{SHAP}| \approx 1.28$ km): the point at which the dispatched pilot boat would meet the ship if both held their course and speed. It is followed by the ship's distance from the breakwater (0.19 km) and the wave height (0.18 km); the remaining features contribute less. The model thus learns a corrected version of the encounter geometry: where the two craft will meet once the boat has been sent. Traffic matters little, consistent with conditions that rarely approach operating limits at this port, while higher waves shift the transfer towards sheltered water. For outbound transfers the leading features are the ship's distance from the breakwater (0.11 km), the pilot boat's bearing from it (0.10 km) and the range between the two (0.07 km), all small against the outbound spread, in line with the weak outbound skill.

5.4. Comparison of CPA metrics and encounter risk

Fig. 7 compares the proximity measures for the subject vessel and its nearest target vessel at the time of closest approach, for the 3168 confirmed transfers with AIS fixes for both ships (heading was replaced by course for 22.9% of target vessels, and default dimensions were used for the 27.3% of targets without AIS dimensions). The centre-to-centre CPA has a median of 1202 m, and 2.3% of encounters come closer than 100 m. Subtracting the half-lengths (Eq. (5)) removes a median of 192 m: the signed separation Δ_{SD} has a median of 1000 m (5th–95th percentile 8–2889 m), is below 100 m in 349 encounters (11.0%) and negative in 149 (4.7%), for which the truncated CPA_{SD} is zero. Point-based CPA thus does not account for ship size, which matters most in the closest encounters.

A negative Δ_{SD} does not show that the hulls overlapped. Of the 149 encounters with $\Delta_{SD} < 0$, the oriented hull rectangles intersect in only 48 (32.2%); the others keep a median polygon clearance of 30 m. Whether these 48 rectangles truly intersected is itself uncertain, because AIS antenna offsets, heading and dimensions are uncertain at the scale of

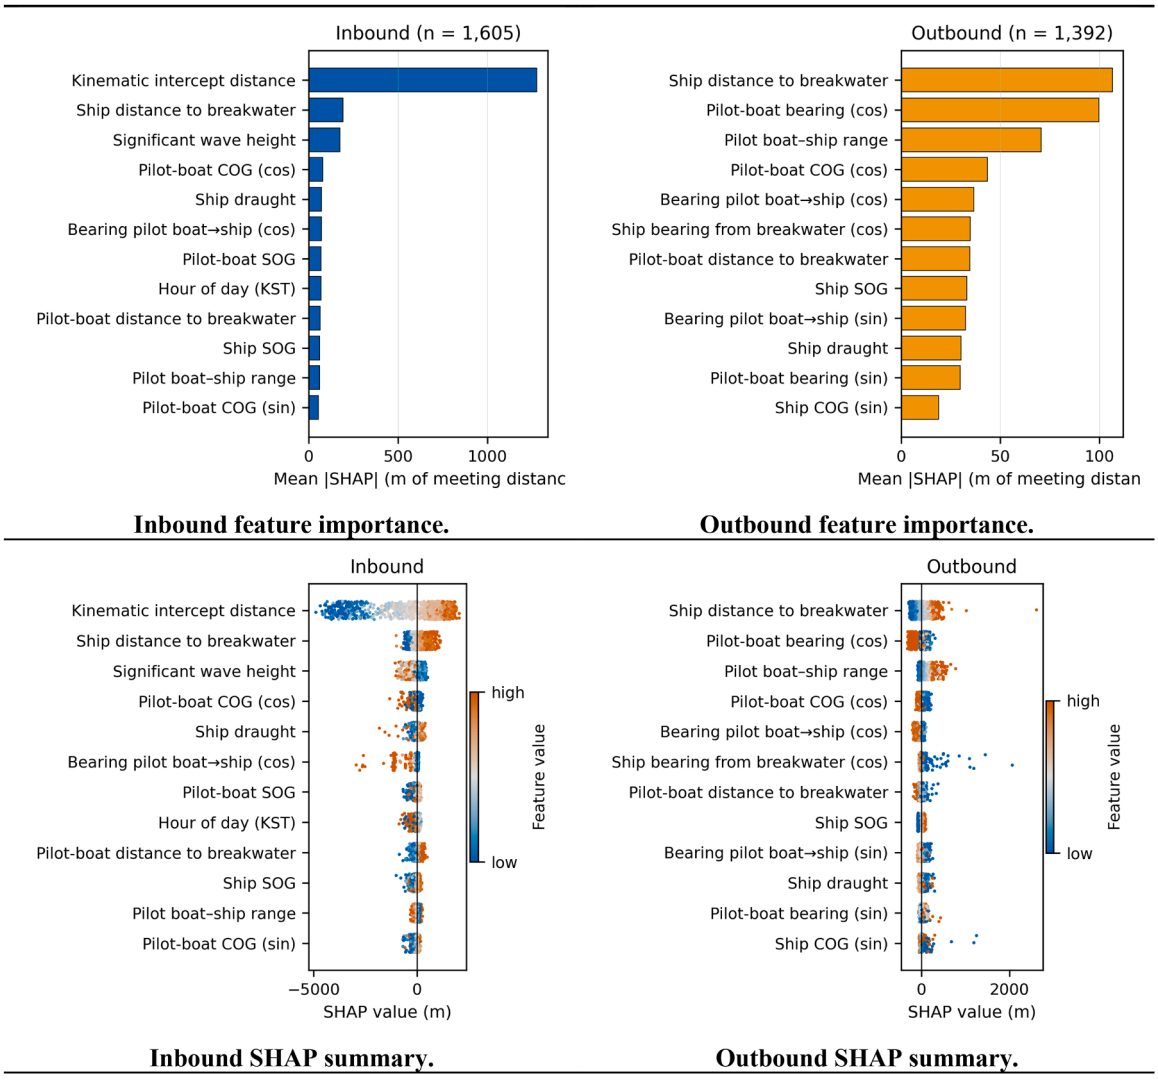

Fig. 6. SHAP feature attributions of the hybrid models.

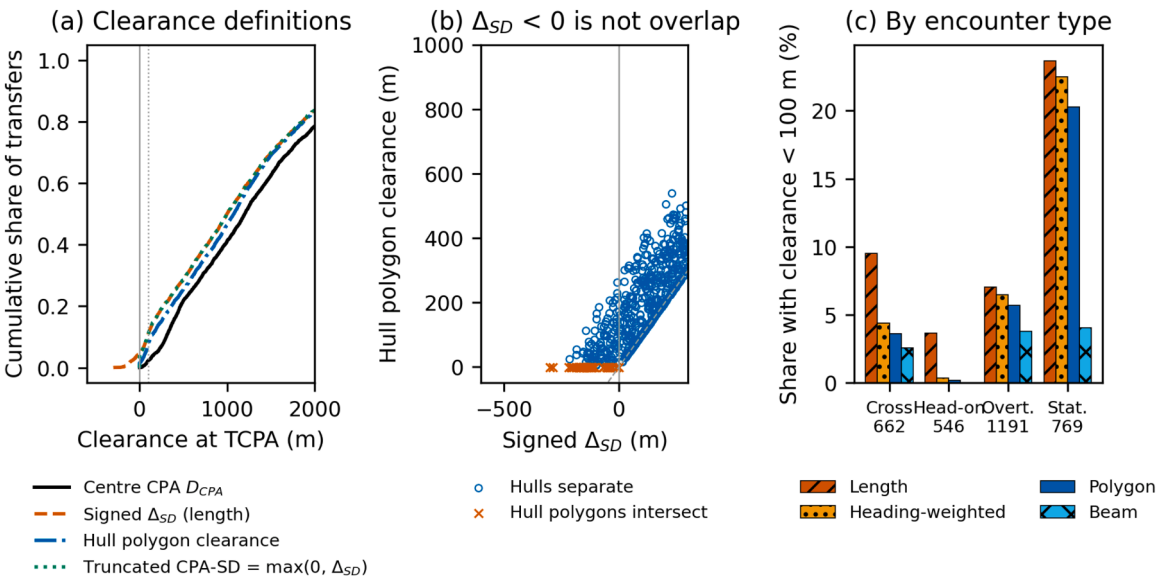

Fig. 7. Proximity at closest approach: (a) Clearance measures; (b) Polygon clearance where $\Delta_{SD} < 0$; (c) share below 100 m.

a few metres.

Table 6 shows that the choice of hull model shifts the clearance distribution and the classification of individual encounters, while a large low-risk majority of recommendations persists under every model (89.4–94.3%). These shares use the hull clearance to the nearest vessel at the decision time; the domain class used in the optimisation subtracts the larger ship-domain radii and gives 83.1%. For the polygon model, the "<0" column counts intersecting rectangles. By encounter type, the length-only model is most conservative where beam is the relevant dimension: the share below 100 m falls from length-only to polygon from 9.5% to 3.6% in crossing encounters ($n = 662$) and from 3.7% to 0.2% head-on ($n = 546$), but only from 7.1% to 5.7% when overtaking ($n = 1191$). Encounters in which one ship moves at <1 kn ($n = 769$) remain the closest under every model (23.7% and 20.3% below 100 m). We retain the length-only CPA_{SD} as the screening metric for two reasons: it needs only the ship lengths, whereas the polygon also needs beam, antenna offsets and heading, which are missing for some AIS targets; and it is the more cautious of the two on our data (median clearance 1000 vs 1077 m; 11.0% vs 7.9% of encounters below 100 m). The polygon clearance is reported alongside it.

5.5. Seasonal and meteorological influences on pilotage dynamics

Wind and wave conditions at Busan Newport are clearly seasonal, strongest in winter (2024Q1: mean wind 7.2 m/s, H_s 1.34 m), but the operational quantities are stable across quarters: the mean meeting distance stays between 3.4 and 4.0 km, and Kruskal–Wallis tests find no significant seasonal difference in meeting distance, Δ_{SD} or session duration, while wind and wave height differ strongly ($p < 10^{-40}$, $\varepsilon^2 \leq 0.084$). Seasonal weather thus changes the conditions of the transfers far more than where and how they take place. The quarterly distributions and tests are given in Supplementary Section S4 (Fig. S4, Tables S2 and S3).

Table 6
Size-corrected clearance under four hull models, and share of feasible recommendations in the low-risk class.

Hull model	Median (m)	<0	<100 m	Low-risk rec.
Length only, Eq. (5)	1000	4.7%	11.0%	89.4%
Beam only	1174	0.3%	2.9%	94.3%
Heading-weighted	1070	3.1%	8.9%	91.9%
Oriented polygon	1077	1.5%	7.9%	92.3%

5.6. Relocation patterns and a worked case

The relocations follow three patterns. (i) *Navigational relocation*: most moves are short shifts of a few hundred metres that bring the predicted point clear of the buoyed approach corridor or into fairway access, or make the water-routed legs of the pilot boat and the ship feasible; these account for most anchors that fail a gate (396 on the buoyed corridor, 387 on navigable water, 137 on fairway access). (ii) *Depth relocation*: 298 anchors have UKC below 2 m, mostly outbound anchors close to the breakwater (211), and are moved to deeper water. (iii) *Traffic relocation*: where the domain clearance to a nearby vessel at dispatch is too small (32 anchors fail the collision gate, and many more lose score in the medium or high class), the point moves, usually along the ship's track, to where the nearest vessel is further away. Because of the relocation penalty, the size of each move reflects how much the safety score gains.

Fig. 8 shows a traffic relocation selected from the results by fixed criteria (inbound, feasible, ship under way at ≥ 5 kn, dispatch at least 10 min before the transfer, anchor failing the collision gate, recommendation 0.5–2.5 km ahead along the ship's course; largest such ship). The inbound container ship MAERSK HALIFAX (LOA 353 m, 151,990 GT) met its pilot at 03:44 KST on 10 May 2024. The pilot boat was dispatched at 03:26 (18.6 min before the transfer), when the ship was 12.1 km out at 5.7 kn and the boat 6.9 km from it. The ML anchor, 8.4 km from the breakwater, fell in the high collision-risk class (nearest vessel 412 m; 1, 1 and 2 vessels within 0.5, 1 and 2 NM). The recommendation lies 2.4 km from the anchor, 2.0 km closer to the port, on the fairway centreline, ahead of the ship and towards the approaching pilot boat, where the collision class is low (nearest vessel 2101 m; 0, 0 and 3 vessels), UKC remains ample (22.9 m against 23.9 m) and the sea state is the same (H_s 1.0 m). The transfer actually took place 1.3 km from the recommendation. The case shows the intended behaviour, and also its limit: the relocation resolves the traffic conflict at the predicted point and moves the meeting onto the fairway, but the prediction itself is about a kilometre off.

5.7. Transferability to other ports

To test whether the decision-point formulation transfers beyond Busan, we applied it to 11 pilot stations with full-year 2023 AIS: six US port approaches (Galveston, Port Arthur, Cape Henry, Southwest Pass, Long Beach and Port Everglades; MarineCadastre data) and five Danish stations (Skagen, Gedser, Drogden, Helsingør and the Great Belt; Danish

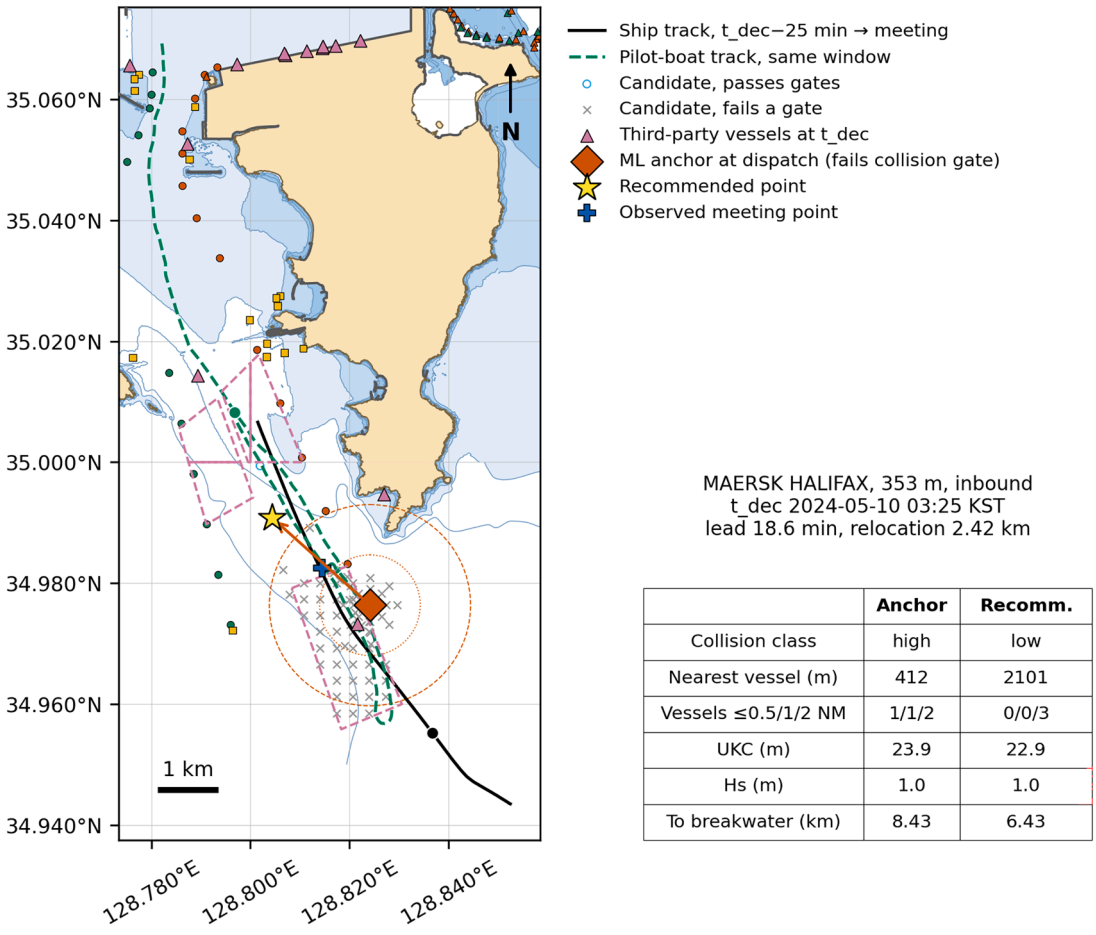

	Anchor	Recomm.
Collision class	high	low
Nearest vessel (m)	412	2101
Vessels ≤0.5/1/2 NM	1/1/2	0/0/3
UKC (m)	23.9	22.9
Hs (m)	1.0	1.0
To breakwater (km)	8.43	6.43

Fig. 8. Worked relocation case: MAERSK HALIFAX, inbound, 10 May 2024.

Maritime Authority data), the latter mostly serving ships in transit through the Danish straits. Pilot boats were identified by AIS ship type 50 and the names of the local pilot fleets, and harbour tugs were excluded. Because no pilotage log was available outside Busan, these transfers are AIS-detected and not log-confirmed. The decision point (pilot-boat dispatch proxy, with the same rule), the predictors and the partitions are those of Busan (Section 4.4, Table 1); wind and wave are ERA5 reanalysis values for the hour containing the dispatch, which is a retrospective product rather than an online forecast. The target is the distance from the transfer to the pilot-boat base, which plays the role of

Table 7
Meeting-distance prediction at 11 further pilot stations, 2023 (MAE in km).

Station	Direction	n	R^2 grouped	MAE	R^2 temporal	B1	Best baseline
Galveston	Inbound	1975	0.566	1.99	0.603	− 2.21	0.559
	Outbound	3174	0.289	2.15	0.308	− 11.0	0.296
Port Arthur	Inbound	1457	0.839	3.21	0.838	− 0.18	0.841
	Outbound	1673	0.713	4.23	0.627	− 1.21	0.712
Cape Henry	Inbound	1458	0.568	0.64	0.286	− 3.27	0.585
	Outbound	1163	0.528	0.56	0.413	− 3.41	0.571
Southwest Pass	Inbound	612	0.593	1.06	0.513	− 6.10	0.597
	Outbound	1450	0.049	0.89	0.025	− 2.00	0.074
Long Beach	Inbound	1373	0.495	1.06	0.460	− 0.48	0.492
	Outbound	2143	0.468	0.46	0.499	− 6.78	0.467
Port Everglades	Inbound	649	− 0.13	1.75	0.353	− 0.91	− 0.14
	Outbound	1294	0.051	0.60	0.136	− 51.0	0.058
Skagen	Inbound	2539	0.737	1.53	0.710	–	0.734
	Outbound	2392	0.447	1.52	0.423	–	0.446
Gedser	Inbound	2839	0.637	0.79	0.556	–	0.634
	Outbound	939	0.832	0.66	0.894	–	0.904
Drogden	Inbound	638	0.755	1.13	0.752	–	0.753
	Outbound	3051	0.776	0.69	0.710	–	0.778
Helsingør	Inbound	2900	0.788	0.92	0.797	–	0.781
	Outbound	908	0.892	1.11	0.854	–	0.891
Great Belt	Inbound	1651	0.921	0.61	0.901	–	0.918
	Outbound	408	0.752	0.85	0.841	–	0.806
Busan (this study)	Inbound	1605	0.495	1.28	0.537	− 0.70	0.494
	Outbound	1392	0.159	0.43	0.027	− 41.8	0.152

the breakwater at Busan. A dispatch was identified for 36,686 transfers, and the median time from dispatch to transfer is 31.0 min inbound and 37.8 min outbound. Unlike at Busan, the AIS coverage includes the berths and pilot bases, and the ship is already visible at dispatch in 85.4% of cases.

Table 7 reports R^2 under ship-grouped 5-fold cross-validation and in a temporal hold-out of November–December, B1 (the official pilot boarding area position, US stations only) and the best of linear regression, random forest and XGBoost. It shows that the dispatch-based prediction carries over to other ports once the model is retrained locally. Inbound R^2 under ship-grouped validation is 0.74–0.92 at four Danish stations (Skagen, Drogden, Helsingør and the Great Belt) and Port Arthur, 0.50–0.64 at Gedser, Cape Henry, Galveston, Southwest Pass and Long Beach, and without skill at Port Everglades ($-$ 0.13, against 0.35 in the temporal hold-out); the temporal hold-out otherwise gives similar values. Outbound prediction is markedly better than at Busan wherever the AIS shows the ship leaving its berth: R^2 is 0.45–0.89 at the five Danish stations, Port Arthur, Cape Henry and Long Beach. It remains weak at Galveston (0.29) and absent at Southwest Pass and Port Everglades (0.05). This suggests that the weak outbound result at Busan is largely a consequence of its AIS extent, which hides the berths and the pilot-boat base. At every US station the official boarding area position predicts the meeting distance worse than the mean. As at Busan, random forest, XGBoost and the hybrid perform within a few hundredths of each other.

Transfer without retraining is unreliable and, at every station, worse than local retraining. Trained on the other stations of the same direction and applied unchanged (leave-one-port-out), the hybrid reaches inbound R^2 of only 0.32 at Skagen and 0.25 at Helsingør, 0.12 or less at the other nine stations, and outbound it has skill only at Helsingør (0.42). The method transfers, but a model fitted at one port must not be used at another without local data and validation.

At the six US stations we also ran the safety-constrained optimisation, with the same causal rules as at Busan (decision at dispatch, past-only AIS, water-routed legs, gates calibrated per cross-validation fold), using official National Oceanic and Atmospheric Administration (NOAA) electronic navigational charts (164 cells; depth areas, dredged areas, soundings, land, lateral marks and charted pilot boarding areas; where a station has no charted boarding area, a fixed 5 km circle around the NOAA boarding position is used) and NOAA tide predictions. A feasible recommendation was found for 77.8–99.8% of the events per station (14,543 of 15,381 in total, 94.6%; lowest at Port Everglades). At Port Everglades all 311 infeasible outbound cases are ships still inside the harbour at dispatch, whose routes must pass the lights inside the narrow entrance channel. The route-clearance gate calibrated there from completed transfers is below 1 m, finer than the 100 m routing grid, so whether a route passes depends on the grid geometry rather than on a real hazard. Point-wise buoy clearance is therefore not a meaningful constraint inside a channel marked by fixed lights, and such ports need channel-aware routing. Of the feasible recommendations, 87.3–93.9% per station are in the low collision-risk class; the median relocation from the anchor is 0–300 m. The UKC constraint, however, could not be applied. US approach channels are maintained by dredging, the charted depth ranges are conservative, and AIS draughts are noisy; with a 2 m rule evaluated at the observed meeting points, only 22% (Port Arthur) and 42% (Galveston) of the completed transfers would pass, and calibrating the rule to observed transfers would require negative thresholds at five stations. We therefore switched the UKC gate off at the US stations and report this as a partial transfer of the constraint set: traffic, buoy, breakwater, fairway and route constraints are evaluated, UKC is not. Applying the UKC constraint at such ports needs surveyed channel depths or reported port depths, one of the local calibrations listed in Section 6.2.

6. Discussion and implications

This study shows that, once the pilot boat has been dispatched, where an inbound transfer will take place can be predicted with moderate accuracy for ships the model has never seen ($R^2 = 0.495$ under ship-grouped validation and 0.537 in later months), far better than the official fixed zones, and that the same holds, often with higher accuracy, at ten of eleven further stations. The decisive information is the encounter geometry at dispatch: where the dispatched boat would intercept the ship on its current course and speed. The earlier the boat is sent, the less this geometry determines the outcome, because the ship's speed and the pilots' choices remain open. Weather and traffic matter little for the prediction, because at these ports conditions rarely approach operating limits; traffic enters the optimisation through the collision gate, whereas wind and wave are recorded for each recommendation but are not enforced as point constraints, because no validated limits exist (Section 4.6). Before dispatch, the meeting point cannot be predicted from AIS: when we used the ship's arrival at a fixed range ring as the decision point, the model had no skill, because where the boat meets the ship depends mainly on when the station sends it.

Outbound transfers at Busan are not predictable from the AIS available for this study. The berths and the pilot-boat base lie outside its extent, the boat is often sent before the departing ship becomes visible, and otherwise only minutes before the disembarkation. At ports where the AIS covers the berths, outbound prediction mostly works (R^2 up to 0.89), so the Busan result reflects the data extent rather than an inherent limit. Operationally, outbound planning at Busan needs information that AIS does not carry, such as scheduled departure times.

Three findings of the submitted version did not survive the revision. First, the two outbound departure regimes were an artefact of harbour-tug interactions in the original fleet list; verified outbound transfers follow a single broad direction. Second, the high accuracy reported earlier (inbound R^2 above 0.9) depended on taking the decision time relative to the observed transfer, which is not possible in operation; at the retrospectively identified dispatch proxy, accuracy is moderate inbound (0.50; 0.30 in the sensitivity without the transfer-time bound) and absent outbound. Third, UKC is not a useful predictor of the meeting distance at this port, because depth rarely constrains boarding; its earlier prominence came from measuring it at the observed meeting point. Its role is that of a hard constraint, where it is active for 15.4% of the predicted points, mostly outbound points close to the breakwater.

The size-corrected separation shows that a centre-to-centre CPA can overstate the clearance between hulls, most of all in the closest encounters, while the hull-polygon check shows that the length-only correction is itself conservative, most strongly in crossing and head-on encounters and least when overtaking (Section 5.4). For screening, the length-only metric is conservative and simple; for incident analysis, the oriented-polygon clearance should be used.

The findings also bear on pilotage policy. IMO Resolution A.960 and national rules define compulsory pilotage areas and nominal boarding grounds but do not distinguish charted "design" locations from operational meeting points. At Busan Newport, transfers take place over a much wider area than the nominal polygons, and the polygon centroid predicts them far worse than any data-driven model (Table 5). A system that recommends points outside the charted areas formalises what pilots already do, but may appear to conflict with a static reading of A.960. The distributions of observed and recommended meeting points could help authorities consider boarding areas as operational corridors and document when transfers inside the breakwater are preferable; such use would require the validation discussed in Section 6.1. The model supports, and does not replace, professional judgement.

6.1. Risk categories as a structured record of the transfer decision

The framework's variables can be mapped onto five recurring categories of pilot-transfer risk. The mapping is conceptual: it records which

factors were evaluated at a proposed point and how, and it is not a validated legal defence, a safe-transfer envelope or an operational certification. Which categories matter, and how they are weighed, depends on the applicable law and the practice of each pilotage authority.

1. Grounding: evaluated by the UKC constraint (Section 4.6): static UKC from charted depth, predicted tide and declared draught at the candidate point, with squat reported as a sensitivity. It is not a route-wide clearance guarantee.
2. Collision with other vessels: evaluated by the domain clearance and traffic counts at the decision time (Section 4.7). The gate threshold is calibrated on completed transfers, traffic is evaluated statically at dispatch, and a feasible recommendation can still fall in the medium or high class.
3. Allision with port infrastructure: evaluated by the breakwater and buoy clearances of the point and of the routed legs. Extending the size-corrected clearance to fixed infrastructure is future work.
4. Occupational injury during ladder transfer: related to the wind and wave conditions of Section 3.2. These are recorded for each recommendation but are not used as point constraints, because no validated limits exist (Section 4.6). Analyses of pilot-transfer accidents identify ladder failure and procedural and environmental factors as the main contributors [14–16].
5. Situational-awareness failure: related to visibility, which is not recorded in the Busan study area, so this category is not operationalised here.

Used in this way, the mapping makes a recommendation and any deviation from it reviewable against the same factors. Its use in regulation or inquiry would require validation against incident data and the applicable legal framework, which is beyond the scope of this study.

6.2. Global applicability and port-level transferability

The input data streams behind the risk categories of Section 6.1 are widely available at the port-approach scale. AIS coverage is mandatory under SOLAS Chapter V for all vessels of 300 GT and above on international voyages; depth is available from official ENCs; tidal data are produced by national oceanographic agencies; meteorological and sea-state data are routinely broadcast by national meteorological services. The framework therefore relies on data classes that are universally available, and not on proprietary or jurisdiction-specific information.

In practical terms, another port needs: (a) AIS coverage that includes the berths and the pilot-boat base, so that the decision time is not limited by the data extent, and in which the pilot boats can be identified by MMSI; (b) a pilotage log or equivalent record to validate the detected transfers, without which non-pilotage contacts such as tug assistance contaminate the data, as they did in the first version of this study; (c) hourly wind and wave observations, ideally from several stations so that inverse-distance weighting can be used; (d) ENC depth, tide predictions and declared draughts; (e) the geometry of the official boarding zones and fairway; (f) at least about one year of transfers; and (g) local recalibration of the breakwater reference, the safety gates and the relocation penalty, and retraining of the models. None of these steps changes the methodology. The evaluation at 11 further stations in the United States and Denmark (Section 5.7) supports this: retrained locally, the dispatch-based model reaches inbound R^2 of 0.50–0.92 at ten of the eleven stations and, where the AIS covers the berths, outbound R^2 up to 0.89, whereas a model transferred unchanged from other ports is unreliable.

6.3. Forward compatibility with maritime autonomous surface ships

The pilot transfer point is a handover zone, where navigational authority changes hands. As MASS regulation develops and remote pilotage emerges [39], the same criteria (traffic, weather window, response time and sea state) could help define where autonomous control yields to shore-based or onboard pilotage, including mixed-traffic approaches and UAV-delivered pilotage, in line with calls for empirically grounded handover boundaries [40,41,50]. The framework does not model cyber-handover latency or remote-operator workload.

7. Conclusions

This study predicts where a pilot transfer will take place at pilot-boat dispatch and relocates the predicted point under hard safety constraints. On 3243 log-confirmed transfers at Busan Newport, inbound meeting distances are predicted with $R^2 = 0.495$ (temporal hold-out 0.537), far better than the official fixed zones, whereas outbound distances are not predictable from the available AIS. The optimisation, using only AIS up to dispatch and chart-based routes, finds a feasible recommendation in 96.4% of cases with a median relocation of 300 m and reports when none exists. The method transfers to other ports when retrained locally. The size-corrected CPA metric, checked against hull polygons, is a conservative screening measure. Together, the results support evidence-based, adaptive pilot transfer corridors and decision support for pilots, VTS and port authorities.

The study has limitations. The AIS data exclude the berths and the pilot-boat base, which limits outbound prediction. The dispatch is a retrospective AIS proxy bounded to 90 min before the transfer, so the results describe retrospective rather than real-time performance. The detector confirms about 28% of the logged transfers, so the data are a validated sample, and the transfers at the other stations are not log-confirmed. Wind and wave come from a single offshore buoy at Busan and from ERA5 reanalysis elsewhere, and visibility is not recorded. Squat is treated as a sensitivity, because vessel-specific hydrodynamic coefficients are not available, and UKC could not be applied at the US ports with charted depths alone. The safety gates are calibrated on completed transfers, and the static evaluation at dispatch does not model traffic movement, so passing the gates is not a proof of safety. Future work will use logged dispatch times, add calibrated prediction intervals to the optimisation, model traffic movement to the meeting time, and extend the size-corrected clearance to fixed infrastructure.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this revision the authors used Claude (Anthropic) and OpenAI Codex to assist with writing and debugging the analysis code (data processing, validation, optimisation and figure scripts), with auditing the code, the reported numbers and the references against the analysis outputs and bibliographic registries, and with language editing of the manuscript and the response to reviewers. All code is released (Data and Code Availability). After using this tool, the authors reviewed and edited the code and the content as needed and take full responsibility for the content of the publication.

CRediT authorship contribution statement

Gokhan Camliyurt: Writing – original draft, Visualization, Software, Methodology, Formal analysis, Data curation, Conceptualization. **Bora San Turgut:** Writing – review & editing, Validation, Methodology. **Hyungoo Park:** Writing – review & editing, Investigation, Data curation. **Sei Hun Kim:** Writing – review & editing, Resources, Investigation. **Daewon Kim:** Writing – review & editing, Investigation, Data curation. **Youngsoo Park:** Writing – review & editing, Supervision, Resources. **Zaili Yang:** Writing – review & editing, Supervision, Project administration, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal

relationships which may be considered as potential competing interests:

Zaili Yang reports financial support was provided by European Research Council (ERC). If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant Agreement No 864724). The authors are grateful to GCS Solution Co., Ltd. (Seoul, Republic of Korea) for compiling and delivering the Korea Meteorological Administration station-level observation records used in this study.

Data availability

The analysis code (transfer detection, log validation, inputs at the decision point, model validation and the safety-constrained optimisation, including the multi-port extension), the list of verified pilot-boat MMSIs and a description of the required inputs are available at https://github.com/gcamliyurt/RESS-pilot-transfer (archived on Zenodo, doi:10.5281/zenodo.23599536). The raw AIS and KMA observation data for Busan were provided under licence and can be shared only with the permission of the data providers; the pilotage log is published by the Busan Harbor Pilots' Association, and the US and Danish AIS and the NOAA charts used for the multi-port analysis are public.

References

[1] IMO. Resolution A.960(23): recommendations on training and certification and operational procedures for maritime pilots other than deep-sea pilots. London: International Maritime Organization; 2003.
[2] IMO. Resolution A.1045(27): pilot transfer arrangements. London: International Maritime Organization; 2011.
[3] IMPA. Information on pilot transfer operations. International Maritime Pilots' Association; 2024.
[4] IALA. IALA VTSManual. 8.5. editor. Saint-Germain-en-Laye: International Association of Marine Aids to Navigation; Lighthouse Authorities; 2025.
[5] Nautical Institute, 2015. Ship hits navigation beacon while disembarking pilot. MARS Report 201572.
[6] Republic of the Marshall Islands Maritime Administrator. Marine safety advisory no. 10-15: safe navigation in the vicinity of designated pilot embarkation /disembarkation areas. Republic of the Marshall Islands Maritime Administrator; 2015.
[7] Rong R, Ma S, Ren N, Lin Q, Jia N. Generative artificial intelligence in intelligent transportation systems: a systematic review of applications. Front Eng Manag 2025;12:1020–36. https://doi.org/10.1007/s42524-025-4241-9.
[8] Li Y, Pu Z, Liu P, Qian T, Hu Q, Zhang J, Wang Y. Efficient predictive control strategy for mitigating the overlap of EV charging demand and residential load based on distributed renewable energy. Renew Energy 2025;240:122154. https://doi.org/10.1016/j.renene.2024.122154.
[9] Chen X, Zhang Z, Li Z, Han B, Zheng Y, Biancardo SA. Microscopic aggregated traffic parameter extraction against complex camera motion interference. Transp Saf Environ 2025;7:tdaf056. https://doi.org/10.1093/tse/tdaf056.
[10] Akyuz E, Celik M. Utilisation of cognitive maps in modelling human error in marine accident analysis and prevention. Saf Sci 2014;70:19–28. https://doi.org/10.1016/j.ssci.2014.05.004.
[11] Camliyurt G, Kang WS, Kim D, Park S, Park Y. A random forest algorithm-based accident prediction to prevent marine pilot occupational accidents. In: Proceedings of the Korean institute of navigation and port research conference; 2022. p. 415–6.
[12] Chauvin C, Lardjane S, Morel G, Clostermann JP, Langard B. Human and organisational factors in maritime accidents: analysis of collisions at sea using the HFACS. Accid Anal Prev 2013;59:26–37. https://doi.org/10.1016/j.aap.2013.05.006.
[13] Camliyurt G, Park Y, Kim D, Kang WS, Park S. Machine learning with multi-source data to predict and explain marine pilot occupational accidents. J Mar Sci Technol 2023;31:2. https://doi.org/10.51400/2709-6998.2709.
[14] Camliyurt G, Choi SA, Kim SR, Guzel AT, Park YS. Risk assessment for marine pilot occupational accidents using fault tree and event tree analysis. J Navig Port Res 2022;46:400–8. https://doi.org/10.5394/KINPR.2022.46.5.400.
[15] Aydın M, Uğurlu Ö, Boran M. Assessment of human error contribution to maritime pilot transfer operation under HFACS-PV and SLIM approach. Ocean Eng 2022;266:112830. https://doi.org/10.1016/j.oceaneng.2022.112830.
[16] Sakar C, Sokukcu M. Dynamic analysis of pilot transfer accidents. Ocean Eng 2023;287:115823. https://doi.org/10.1016/j.oceaneng.2023.115823.
[17] Butler GL, Read GJM, Salmon PM. Understanding the systemic influences on maritime pilot decision-making. Appl Ergon 2022;104:103827. https://doi.org/10.1016/j.apergo.2022.103827.
[18] Oraith H, Blanco-Davis E, Yang Z, Matellini B. An evaluation of the effects of human factors on pilotage operations safety. J Mar Sci Appl 2021;20:393–409. https://doi.org/10.1007/s11804-021-00222-1.
[19] Nautical Institute, 2009. Collision near pilot station. MARS Report 200945.
[20] Nautical Institute, 1994. Congestion at pilot station – Wandelaar. MARS report 94007.
[21] Zhang W, Goerlandt F, Montewka J, Kujala P. A method for detecting possible near miss ship collisions from AIS data. Ocean Eng 2015;107:60-9–9. https://doi.org/10.1016/j.oceaneng.2015.07.046.
[22] Goerlandt F, Kujala P. Traffic simulation based ship collision probability modeling. Reliab Eng Syst Saf 2011;96:91–107. https://doi.org/10.1016/j.ress.2010.09.003.
[23] Rong H, Teixeira AP, Guedes Soares C. Data mining approach to shipping route characterization and anomaly detection based on AIS data. Ocean Eng 2020;198:106936. https://doi.org/10.1016/j.oceaneng.2020.106936.
[24] Huang L, Wen Y, Zhang Y, Zhou C. Dynamic calculation of ship exhaust emissions based on real-time AIS data. Transp Res D: Transp Environ 2020;80:102277. https://doi.org/10.1016/j.trd.2020.102277.
[25] Chen X, Zhang Y, Feng Z, Zhang H, Biancardo SA, Shang WL. Spatiotemporal ship carbon emission exploitation and low-carbon shipping pathway evaluation. Appl Energy 2026;422:128334. https://doi.org/10.1016/j.apenergy.2026.128334.
[26] Chen X, Liu Y, Achuthan K. A semi-supervised deep learning model for ship encounter situation classification. Ocean Eng 2021;239:109824. https://doi.org/10.1016/j.oceaneng.2021.109824.
[27] Chen X, Wu P, Wang Z, Feng Z, Luo L, Zhang H, Biancardo SA. MKG-GNN: maritime knowledge graph and GNN framework for ship speed forecasting in port. Ocean Eng 2026;355:125179. https://doi.org/10.1016/j.oceaneng.2026.125179.
[28] Korupoju AK. Ship collision risk evaluation using AIS and weather data through deep learning. Ocean Eng 2025;318:120116. https://doi.org/10.1016/j.oceaneng.2024.120116.
[29] Munim ZH, Sørli MA, Kim H, Alon I. Predicting maritime accident risk using automated machine learning. Reliab Eng Syst Saf 2024;248:110148. https://doi.org/10.1016/j.ress.2024.110148.
[30] Murray B, Perera LP. An AIS-based deep learning framework for regional ship behavior prediction. Reliab Eng Syst Saf 2021;215:107819. https://doi.org/10.1016/j.ress.2021.107819.
[31] Gucma L, Bilewski M, Artyszuk J, Drwiega K. System for determining dynamic under keel clearance of vessels entering the port of Swinoujscie (DRWPS). TransNav Int J Mar Navig Saf Sea Transp 2017;11:539–45. https://doi.org/10.12716/1001.11.03.21.
[32] Wu L, Jia S, Wang S. Pilotage planning in seaports. Eur J Oper Res 2020;287:90–105. https://doi.org/10.1016/j.ejor.2020.05.009.
[33] Cafieri S, Conn AR, Mongeau M. Mixed-integer nonlinear and continuous optimization formulations for aircraft conflict avoidance via heading and speed deviations. Eur J Oper Res 2023;310:670–9. https://doi.org/10.1016/j.ejor.2023.03.002.
[34] Peyronne C, Conn AR, Mongeau M, Delahaye D. Solving air traffic conflict problems via local continuous optimization. Eur J Oper Res 2015;241:502–12. https://doi.org/10.1016/j.ejor.2014.08.045.
[35] Dell'Olmo P, Lulli G. A dynamic programming approach for the airport capacity allocation problem. IMA J Manage Math 2003;14:235–49. https://doi.org/10.1093/imaman/14.3.235.
[36] Eser E, Monteil J, Simonetto A. On the tracking of dynamical optimal meeting points. IFAC-PapersOnLine 2018;51:434–9. https://doi.org/10.1016/j.ifacol.2018.07.071.
[37] Liu K, Liu J, Zhang J. Advanced algorithms for optimal meeting points in road networks. IET Intell Transp Syst 2023;17:1022–36. https://doi.org/10.1049/itr2.12323.
[38] Otaki K. Planning with explanations for finding desired meeting points on graphs. In: Proceedings of the AAAI conference on artificial intelligence; 2022. p. 10319–26. https://doi.org/10.1609/aaai.v36i9.21273.
[39] Kurt I, Aymelek M. Operational adaptation of ports with maritime autonomous surface ships. Transp Policy (Oxf) 2024;145:1–10. https://doi.org/10.1016/j.tranpol.2023.09.023.
[40] Li X, Chua JY, Yuen KF. A review on maritime disruption management: categories, impacts, and strategies. Transp Policy (Oxf) 2024;154:40–7. https://doi.org/10.1016/j.tranpol.2024.05.013.
[41] Li X, Oh P, Zhou Y, Yuen KF. Operational risk identification of maritime surface autonomous ship: a network analysis approach. Transp Policy (Oxf) 2023;130:1–14. https://doi.org/10.1016/j.tranpol.2022.10.012.
[42] GEBCO Bathymetric Compilation Group 2024, 2024. The GEBCO_2024 grid – a continuous terrain model of the global oceans and land. https://doi.org/10.5285/1c44ce99-0a0d-5f4f-e063-7086abc0ea0f.
[43] Mou JM, van derTak C, Ligteringen H. Study on collision avoidance in busy waterways by using AIS data. Ocean Eng 2010;37(5-6):483–90. https://doi.org/10.1016/j.oceaneng.2010.01.012.
[44] Sinnott RW. Virtues of the haversine. Sky Telescope 1984;68:159.
[45] Busan Harbor Pilots' Association, 2024. Passage plan for Busan new port.
[46] Chai T, Draxler R. Root mean square error (RMSE) or mean absolute error (MAE)? – arguments against avoiding RMSE in the literature. Geosci Model Dev 2014;7:1247–50. https://doi.org/10.5194/gmd-7-1247-2014.
[47] Draper N, Smith H. Applied regression analysis. Wiley; 1998.

[48] Willmott CJ, Matsuura K. Advantages of the mean absolute error (MAE) over the root mean square error (RMSE). Clim Res 2005;30:79–82. https://doi.org/10.3354/cr030079.
[49] Lundberg S, Lee SI. A unified approach to interpreting model predictions. Advances in neural information processing systems. Curran Associates Inc.; 2017. p. 4765–74.
[50] Yap WY, Hsieh CH, Lee PTW. Shipping connectivity data analytics: implications for maritime policy. Transp Policy (Oxf) 2023;132:112–27. https://doi.org/10.1016/j.tranpol.2022.12.018.