Supplementary Materials

Feature Transformer and LightGBM Ensemble for Multi-class Ship Trajectory Recognition Using Real AIS Data

This document provides supplementary figures and tables referenced in the main manuscript. It contains four supporting figures (S1–S4) and two tables (S1–S2). All experiments were performed on the same 70/15/15 trajectory-level split of 2196 trajectories from 14 ship types described in Section 4 of the main paper.

1. Figure S1 — Training and validation curves

Figure S1 shows the training loss together with the training and validation accuracy of the Feature Transformer over its pre-training epochs. The model exhibits stable convergence: the training loss decreases monotonically, while the validation accuracy plateaus, indicating that the early-stopping criterion (patience = 10) is conservative for this dataset size.

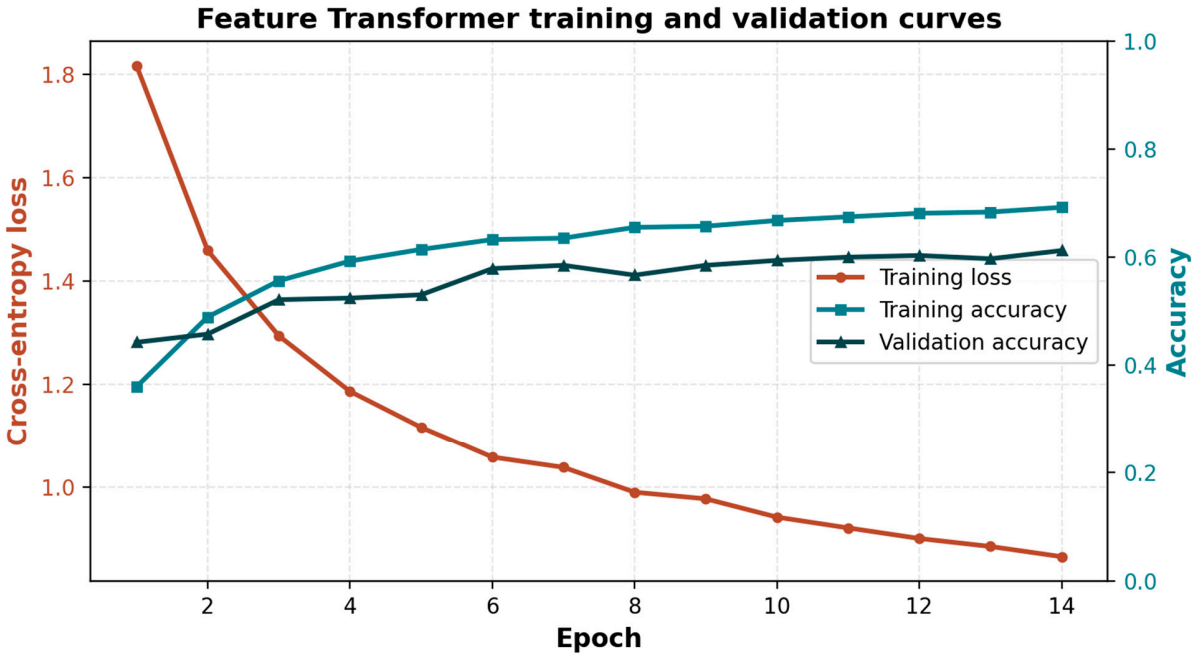


Figure S1. Feature Transformer training loss (left axis) together with training and validation accuracy (right axis) on the AIS classification task.

2. Figure S2 — Per-class precision, recall, and F1

Figure S2 compares per-class precision, recall, and F1 between the baseline Hybrid model and the same model with SMOTE oversampling applied during training. Classes are ordered by support (ascending) to

emphasise the long-tailed distribution. SMOTE consistently improves recall on minority classes (LNG Carrier, Pleasure Craft, Other Cargo) while leaving majority-class metrics nearly unchanged.

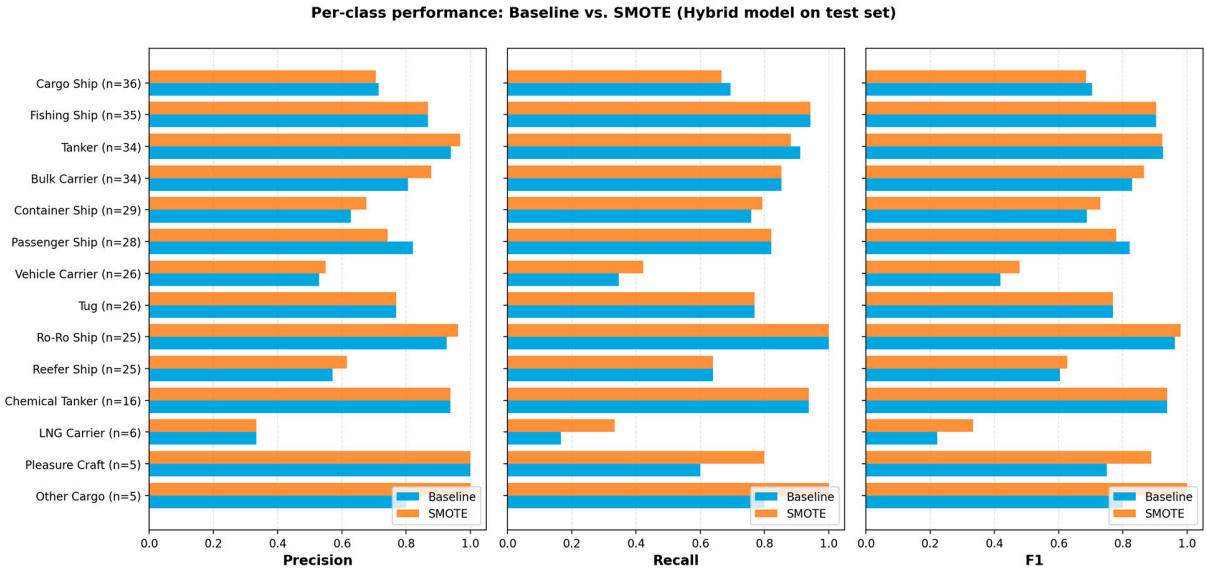


Figure S2. Per-class precision, recall, and F1 of the Hybrid model on the test split, comparing baseline training and SMOTE oversampling. Class names are followed by their test support n.

3. Figure S3 — Top-5 SHAP features per class

Figure S3 ranks the five features with the largest mean absolute SHAP value for each of the 14 ship types, computed on the test split using the LightGBM TreeExplainer. The plots show that Transformer embedding dimensions (trans_emb_*) dominate the explanation for most classes, while a small number of hand-crafted tabular features (duration, traj_len, dt_max, stop_time, curv_max, sog_range, cog_max) play distinctive roles for specific ship types such as Bulk Carrier, Pleasure Craft, and Tanker.

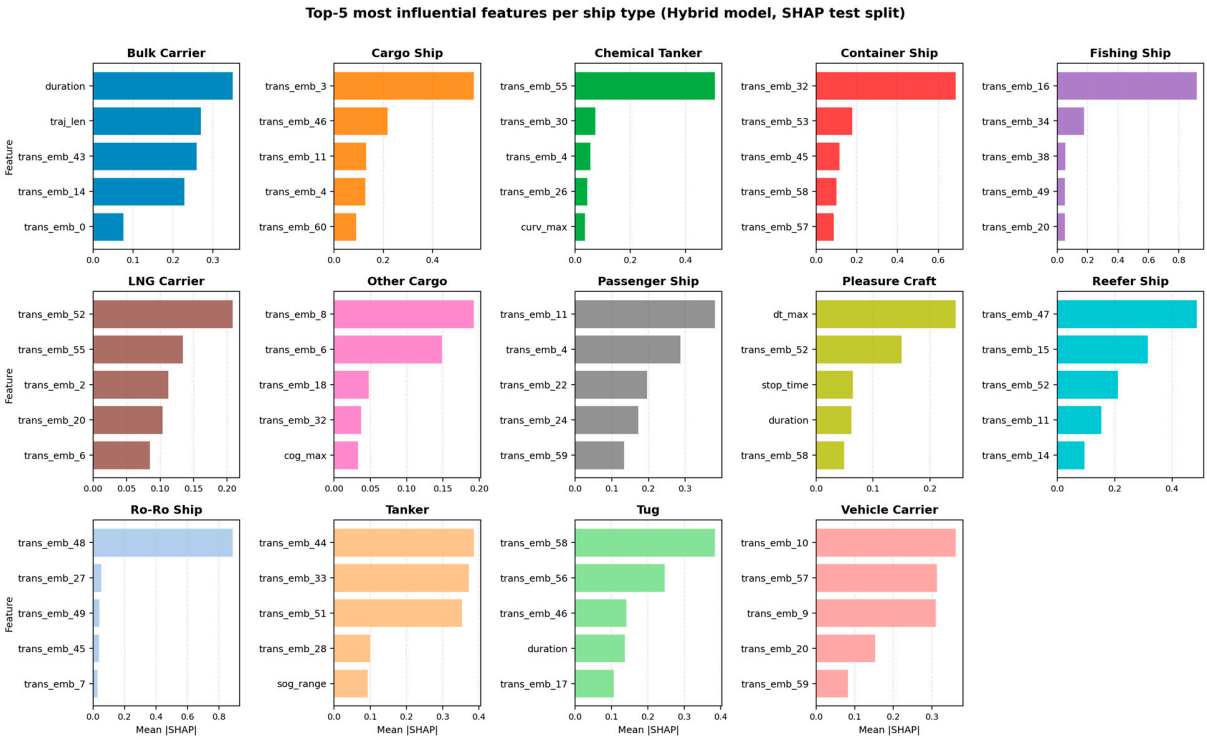


Figure S3. Top-5 features ranked by mean absolute SHAP value, plotted separately for each of the 14 ship types. Feature names starting with trans_emb_ denote individual Transformer fingerprint dimensions; the remaining names refer to hand-crafted tabular features.

5. Table S1 — Hyperparameter settings

Table S1 reports the full set of hyperparameters used to obtain the main results in the manuscript, including (i) the Feature Transformer architecture and training schedule, (ii) the LightGBM hyperparameters (both manual defaults and the Optuna-tuned values), and (iii) the Optuna search space.

Table S1. Hyperparameter settings for the Feature Transformer, LightGBM, and the Optuna search space.

Component	Hyperparameter	Value / Search range
Feature Transformer	Input dimension F	4 (lon, lat, SOG, COG)
	Sequence length L (after sliding-window)	128
	Number of segments per trajectory	3 (first / middle / last)
	Segment aggregation	Mean of segment-level fingerprints
	Embedding dimension d_model	64 (default; sweep over 32/64/128/256 in Section 5.3.3)
	Encoder layers	2
	Attention heads	4
	FFN dimension	256
	Positional encoding	Sinusoidal, max_len = 128

	Dropout	0.1
	Pooling	Last-token pooling
	Optimizer	Adam
	Learning rate	1×10^{-3}
	Batch size	64
	Max epochs	100 (early stopping patience = 10 on validation accuracy)
LightGBM (Optuna-tuned)	num_leaves	24
	max_depth	6
	learning_rate	0.0241
	min_data_in_leaf	36
	feature_fraction	0.584
	bagging_fraction	0.950
	bagging_freq	6
	lambda_l1	1.59×10^{-3}
	lambda_l2	9.78×10^{-4}
	Boosting rounds	500 (early stopping = 50)
Optuna search space	num_leaves	Int [16, 96]
	max_depth	Int [4, 14]
	learning_rate	Float [0.01, 0.15], log-uniform
	min_data_in_leaf	Int [5, 80]
	feature_fraction	Float [0.5, 1.0]
	bagging_fraction	Float [0.5, 1.0]
	bagging_freq	Int [0, 7]
	lambda_l1	Float [1×10^{-8}, 1.0], log-uniform
	lambda_l2	Float [1×10^{-8}, 1.0], log-uniform
	Number of trials	50
	Optimization target	Validation Macro-F1

6. Table S2 — Full list of 103 hand-crafted tabular features

Table S2 (see file Table_S2_features.docx) enumerates the 103 hand-crafted features used by LightGBM. Features are organised into five groups: (i) kinematic statistics on SOG and COG, (ii) geometric and shape descriptors of the trajectory, (iii) general statistical moments of lon/lat/SOG/COG, (iv) frequency-domain features (low-order FFT coefficients of SOG and COG), and (v) segment-level summaries.