Computers, Materials & Continua

Tech Science Press

https://doi.org/10.32604/cmc.2026.077222

ARTICLE

Secure and Differentially Private Edge-Cloud Federated Learning Framework for Privacy-Preserving Maritime AIS Intelligence

Abuzar Khan[1], **Abid Iqbal**[2,*], **Ghassan Husnain**[1,*], **Fahad Masood**[1], **Mohammed Al-Naeem**[3] and **Sajid Iqbal**[4]

[1]Department of Computer Science, CECOS University of IT and Emerging Sciences, Peshawar, Pakistan
[2]Department of Computer Engineering, College of Computer Sciences and Information Technology, King Faisal University, Al-Ahsa, Saudi Arabia
[3]Department of Computer Networks Communications, CCSIT, King Faisal University, Al Ahsa, Saudi Arabia
[4]Department of Information Systems, College of Computer Science and Information Technology, King Faisal University, Al-Ahsa, Saudi Arabia

*Corresponding Authors: Abid Iqbal. Email: aaiqbal@kfu.edu.sa; Ghassan Husnain. Email: ghassan.husnain@cecos.edu.pk

Received: 04 December 2025; Accepted: 19 January 2026; Published: 09 April 2026

ABSTRACT: Cloud computing now supports large-scale maritime analytics, yet offloading rich Automatic Identification System (AIS) data to the cloud exposes sensitive operational patterns and complicates compliance with cross-border privacy regulations. This work addresses the gap between growing demand for AI-driven vessel intelligence and the limited availability of practical, privacy-preserving cloud solutions. We introduce a privacy-by-design edge-cloud framework in which ports and vessels serve as federated clients, training vessel-type classifiers on local AIS trajectories while transmitting only clipped, Gaussian-perturbed updates to a zero-trust cloud coordinator employing secure and robust aggregation. Using a public AIS corpus with realistic non-IID client partitions, our evaluation shows that non-private FedAvg attains validation AUC ≈ 0.90 and test AUC ≈ 0.78, closely matching a centralized baseline. Moderate differential privacy noise ($\delta \leq 0.5$) preserves most of this utility across KRUM and trimmed-mean aggregation. Communication analysis indicates that secure aggregation introduces negligible overhead compared with standard FedAvg, while homomorphic encryption increases payload size by roughly an order of magnitude. Membership-inference experiments further demonstrate strong privacy protection, yielding ROC AUC ≈ 0.51 with no correctly inferred training members. Overall, the findings show that effective, regulation-conscious maritime analytics can be achieved without centralizing raw AIS data, offering a practical pathway for deploying resilient, privacy-enhanced AI services in distributed maritime environments.

KEYWORDS: Federated learning; privacy-preserving cloud computing; maritime AIS analytics; differential privacy

1 Introduction

Cloud computing now underpins large-scale analytics in safety-critical domains, including maritime logistics [1]. Modern vessels stream AIS data to the cloud to enable AI-driven services such as route optimisation, traffic forecasting, collision avoidance and anomaly detection [2]. However, centralising raw AIS data from diverse stakeholders raises major security and privacy risks: trajectories can expose commercial strategies and operational weaknesses and cross-border data sharing must satisfy varied regulations [3]. Perimeter security and basic anonymisation are inadequate, as AIS trajectories remain highly re-identifiable and cloud insiders can still infer sensitive patterns [4].

High-quality maritime AI requires diverse data across vessel types, regions and conditions [5]. Yet operators cannot always share raw AIS feeds and isolated local models lack the robustness needed for reliable cloud-based decision support [5]. These tensions highlight the need for privacy-preserving, collaborative edge-cloud analytics in maritime infrastructures [6]. Fig. 1 shows the heterogeneous, distributed environment that motivates applying federated learning in this domain [7].

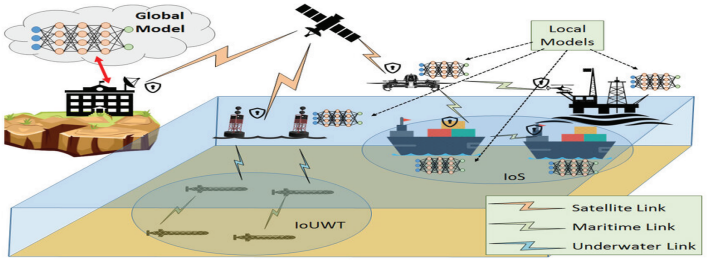

Figure 1: Federated learning in various maritime scenarios, including Internet of Ships (IoS), maritime IoT with buoys and Internet of Underwater Things (IoUWT), adapted from Giannopoulos et al. [7].

We introduce a privacy-preserving edge-to-cloud collaborative learning framework for AIS-based maritime analytics [8]. In our design, vessels and port infrastructures act as federated clients, retaining raw AIS data locally and transmitting only clipped, noise-perturbed model updates instead of original records [9]. Privacy is enforced at the learning layer through differential privacy, which bounds information leakage from individual trajectories [10], while secure aggregation masks updates so the server observes only jointly aggregated parameters rather than client-specific contributions [11]. This algorithmic protection achieves privacy without depending on external security services. Using a realistic Kaggle AIS dataset, we evaluate performance under non-IID partitions and multiple threat models, considering utility, communication cost, and inference-attack resistance [12]. Results show that anomaly detection and route modelling remain feasible under formal privacy guarantees, supporting practical maritime deployment. The framework therefore provides a reproducible benchmark for federated learning and privacy-enhancing technologies in real-world cyber-physical systems [13].

2 Related Work

Research on cloud security, PETs and distributed learning has grown rapidly [14]. Two relevant strands are: (i) privacy-preserving and federated learning for cloud-edge systems and (ii) AIS-based maritime analytics in security-critical settings [15]. Taken together, they reveal a gap between generic privacy-preserving AI methods and domain-specific maritime needs motivating our integrated, privacy-aware edge-to-cloud framework.

2.1 Privacy Preserving and Federated Learning in Cloud and Edge Environments

Federated learning (FL) enables distributed training when data cannot be centrally stored due to confidentiality, regulation, or operational constraints [16]. Early work demonstrated that simple model averaging remains viable despite heterogeneous devices, intermittent participation, and bandwidth limitations, while later research adapted FL to edge–cloud settings to handle non-IID data, unbalanced client sizes, and straggler effects [17]. To strengthen confidentiality, privacy-enhancing technologies such as differential privacy and secure aggregation have been incorporated, supporting zero-trust deployment scenarios [18]. However, much of the literature still relies on simplified benchmarks and overlooks the interoperability,

privacy, and safety demands of critical infrastructures, often emphasizing accuracy while downplaying trade-offs involving privacy budgets, model fidelity and communication overhead [19]. Our work addresses this gap using realistic spatio-temporal data and threat-aware evaluation.

2.2 Maritime Analytics and AIS Based Security Applications

AIS data underpins maritime situational awareness, enabling traffic characterisation, vessel behaviour modelling, route prediction, anomaly detection and collision avoidance [20]. Machine and deep learning methods have been applied to ship type classification, ETA prediction and abnormal behaviour detection, often assuming access to large central AIS repositories operated by authorities or commercial providers [21]. While effective, this centralised paradigm overlooks confidentiality, cross-border data regulations and sensitivities around vessel operations [20].

Maritime cybersecurity research has also exposed AIS vulnerabilities including spoofing, jamming and inference attacks [22]. Existing defences mainly rely on message-level cryptography or coarse anonymisation [23,24], which do not support fine-grained learning tasks or provide formal privacy guarantees [7]. To our knowledge, no prior work combines federated learning, differential privacy and secure aggregation for AIS analytics in edge–cloud maritime environments, despite early FL demonstrations in related tasks [7]. By treating vessels and ports as clients and evaluating on a realistic AIS corpus, we address this gap. Table 1 summarizes prior work, limitations and our contributions.

Table 1: Related work, gaps and how our framework addresses them.

Ref.	Approach	Key Limitation	Our Contribution
[14]	FL and PET surveys	Limited real non-IID, safety-critical evaluation	Full FL + DP + SecureAgg on AIS data
[16,25]	FL optimisation for edge/cloud	Uses generic benchmarks only	Domain-aware FL for maritime settings
[26,27]	DP and secure aggregation	Not applied to spatio-temporal tasks	Integrated DP + SecureAgg for AIS analytics
[20]	Centralised maritime ML	Ignores privacy and data-sovereignty limits	Federated, privacy-preserving AIS modelling
[22]	AIS security and threat analysis	No full learning framework with privacy	Robust FL against inference/poisoning
[7]	Early maritime FL demos	Lacks DP, SecureAgg and benchmarks	Unified reproducible pipeline with threats

3 Methodology

This section outlines our end-to-end privacy-preserving federated AIS analytics framework, covering data preprocessing, client partitioning, model training and attack-defense evaluation.

3.1 Dataset

We use a public AIS trajectory dataset from Kaggle [28], consisting of time-stamped maritime broadcasts with vessel kinematics, identifiers and operational metadata. After removing incomplete entries, timestamp errors, and speed outliers, the final corpus contains approximately ~1.2 M AIS messages from more than 5000 unique vessels operating across multiple ports in the Kattegat Strait region. Each entry

includes spatial coordinates, navigation state and ship characteristics, which collectively define the movement behaviour of individual vessels. To support federated learning, we add dynamic features (speed, heading change, region tags) and assign a `client_id` per vessel or region to create realistic non-IID client partitions for evaluation.

3.2 Data Preprocessing and Feature Engineering

Before training, we preprocess the AIS data in Table 2 to form clean, time-ordered vessel trajectories. For each `mmsi`, messages are sorted and invalid timestamps removed. To ensure data validity, we apply rule-based cleaning to discard corrupted or inconsistent records: (i) non-monotonic or reversed timestamps ($t_i \leq t_{i-1}$), (ii) duplicate timestamps or parsing failures, (iii) abnormal temporal gaps ($\Delta t > 1$ h) indicating corruption, (iv) timezone-inconsistent entries that break chronological ordering, (v) out-of-range coordinates beyond maritime bounds, and (vi) speed anomalies where implied velocity exceeds physical feasibility ($v_i >$ 50 kn) or contradicts reported `sog`. Short gaps are linearly interpolated to maintain continuous trajectories.

Table 2: Core AIS dataset attributes used in our FL experiments (spatio-temporal context: `timestamp`, `lat`, `lon`).

Attr.	Example	Attr.	Example	Attr.	Example
`timestamp`	2019-06-01 12:34:56	`Sog`	12.3 kn	`Region`	North_Sea
`mmsi`	244660123	`Cog`	182.5 deg	`client_id`	client_v1
`lat`	52.41	`Heading`	90 deg	`anomaly_label`	0/1
`lon`	4.87	`ship_type`	Cargo	`destination`	ROTTERDAM

Note: The combination of `timestamp` and `lat`/`lon` defines the vessel's spatio-temporal trajectory. These patterns are central to anomaly detection and enable realistic non-IID client structures when grouped by vessel or port for federated learning.

Given a trajectory $\{(\varphi_i, \lambda_i, t_i)\}_{i=1}^{N}$, we compute the inter-sample time difference

$$\Delta t_i = t_i - t_{i-1}, \qquad i = 2, \ldots, N, \tag{1}$$

and discard segments with large gaps, while small gaps are linearly interpolated. Spatial displacement is computed using the great-circle (Haversine) distance d_i defined as

$$d_i = 2R \arcsin\left(\sqrt{\sin^2\left(\frac{\varphi_i - \varphi_{i-1}}{2}\right) + \cos(\varphi_{i-1})\cos(\varphi_i)\sin^2\left(\frac{\lambda_i - \lambda_{i-1}}{2}\right)}\right), \tag{2}$$

where R is the Earth radius. Based on Eqs. (1) and (2), we derive the instantaneous speed over ground v_i as

$$v_i = \frac{d_i}{\Delta t_i}, \tag{3}$$

and convert it to knots for consistency with `sog`. Unrealistic values (e.g., $v_i > 50$ knots) are removed as outliers. We also derive higher-level features such as heading change in Eq. (4):

$$\Delta\theta_i = \mathrm{wrap}\big(\theta_i - \theta_{i-1}\big), \tag{4}$$

where θ_i is the vessel heading or course and wrap maps angles to $(-\pi, \pi]$. Region labels come from clustering (φ_i, λ_i) and `client_id` follows vessel or port grouping (Table 3). All features are standardized per client before model training in Section 3.5.

Table 3: Engineered features derived from raw AIS data (Table 2).

Feature	Description	From	Feature	Description	From
Δt_i	Time gap between messages	`Timestamp`	$\Delta\theta_i$	Heading/course change	`Cog`, `heading`
d_i	Great-circle distance	`Lat`, `lon`	`Region`	Spatial cluster label	`lat`, `lon`
v_i	Instantaneous speed	d_i, Δt_i	`client_id`	FL client assignment	`mmsi`, `region`

3.3 Federated Client Partitioning and Privacy-Preserving Training

We simulate a practical edge–cloud FL environment in which a central coordinator (cloud node) orchestrates 20–50 virtual maritime clients per round, each emulating either a vessel or port. The server runs on a single cloud instance (Ubuntu 20.04, 8 vCPUs, 32 GB RAM) with a standard PyTorch-based federated loop, and clients communicate with the coordinator over secure gRPC-style channels. Model updates are clipped/noised according to the differential privacy parameters in Table 4. Secure Aggregation is applied prior to averaging, ensuring that only masked model updates are visible to the coordinator.

Table 4: Edge–cloud FL simulation configuration and client partitioning.

Category	Configuration	Notes
Server role	Cloud node Ubuntu 20.04, 8 vCPU, 32 GB RAM	Central coordinator in edge–cloud setup
Client count	20–50 virtual vessels or ports per round	Each client holds its own data shard
Communication	Secure gRPC-style channels	No raw AIS data transmitted
Local training	2–5 epochs per round, batch size 128	Constrained by client resources and per-round time budget
Privacy layer	DP clipping, Gaussian noise, Secure Aggregation	Only masked updates reach the server
IID scheme	Random vessel groups	Nearly balanced distributions across clients
Label skew	Dirichlet-based label splits	Controlled imbalance by alpha value
Region aware	Clients mapped to ports or regions	Natural heterogeneity from geography

From this setup, virtual clients are instantiated from the cleaned AIS corpus $\mathcal{D}$ (Section 3.2). Data are grouped by vessel `mmsi` or region `region` (Table 4), forming disjoint client subsets as defined in Eq. (5):

$$\mathcal{D} = \bigcup_{k=1}^{K} \mathcal{D}_k, \qquad \mathcal{D}_i \cap \mathcal{D}_j = \varnothing \ (i \neq j). \tag{5}$$

Each client further splits its data temporally into training, validation and test sets in Eq. (6):

$$\mathcal{D}_k = \mathcal{D}_k^{\text{train}} \cup \mathcal{D}_k^{\text{val}} \cup \mathcal{D}_k^{\text{test}}. \tag{6}$$

We vary heterogeneity using the three schemes in Table 4. The IID vessel-based scheme randomizes vessel assignment, while the label-skewed scheme samples client label weights from a Dirichlet prior,

$$\boldsymbol{\pi}^{(k)} \sim \text{Dirichlet}(\alpha \mathbf{1}), \tag{7}$$

where smaller α yields stronger skew. Region-aware clients preserve natural traffic imbalance. Client drift is quantified via KL divergence in Eq. (8):

$$\text{KL}(p_k \,\|\, p) = \sum_y p_k(y) \log \frac{p_k(y)}{p(y)}. \tag{8}$$

Training follows a federated protocol using the global objective in Eq. (9):

$$\min_{\mathbf{w}} F(\mathbf{w}) = \sum_{k=1}^{K} \frac{n_k}{n} F_k(\mathbf{w}), \qquad F_k(\mathbf{w}) = \frac{1}{n_k} \sum_{(x,y)\in\mathcal{D}_k} \ell(f_{\mathbf{w}}(x), y). \tag{9}$$

Clients compute local updates $\mathbf{w}_t^k$ and the server aggregates via FedAvg in Eq. (10):

$$\mathbf{w}_{t+1} = \sum_{k\in\mathcal{S}_t} \frac{n_k}{\sum_{j\in\mathcal{S}_t} n_j} \mathbf{w}_t^{(k)}. \tag{10}$$

To ensure privacy, we adopt DP-FedAvg by clipping client updates in Eq. (11)

$$\tilde{\Delta\mathbf{w}}_t^{(k)} = \Delta\mathbf{w}_t^{(k)} \cdot \min\left(1, \frac{C}{\|\Delta\mathbf{w}_t^{(k)}\|_2}\right), \tag{11}$$

and adding Gaussian noise,

$$\hat{\Delta\mathbf{w}}_t^{(k)} = \tilde{\Delta\mathbf{w}}_t^{(k)} + \mathcal{N}(\mathbf{0}, \sigma^2 C^2 \mathbf{I}), \tag{12}$$

yielding (ϵ, δ) privacy (Section 4). Secure aggregation hides individual updates via cryptographic masking before applying Eq. (10). Privacy-related configurations appear in Table 5.

Table 5: Privacy-preserving FL configurations for Section 3.3.

Config	DP Mechanism	Target Privacy	SecureAgg
FedAvg baseline	None	N/A	No
Local DP-FedAvg	Client clipping + Gaussian noise	Moderate (ϵ, δ)	Optional
Central DP-FedAvg	Noise on aggregated update	Strong per-sample privacy	No
DP-FedAvg + SecureAgg	Clipping + noise + masking	Same DP with server protection	Yes

3.4 Attack Models and Defense Mechanisms

To assess robustness and privacy in Section 3.3, we examine standard inference and poisoning attacks in the federated AIS threat model. Eq. (13) reports the membership inference attack success rate, which quantifies privacy leakage by measuring how well an attacker can infer whether a sample (x, y) participated in training. In contrast, Eq. (14) (backdoor poisoning success) measures model integrity compromise by

quantifying how effectively an attacker can manipulate predictions via poisoned updates. Together, these metrics capture two complementary attacker goals in FL for AIS: inferring training participation (privacy) vs. forcing targeted misclassification (integrity), motivating robust defence mechanisms.

$$\mathrm{ASR}_{\mathrm{MI}} = \Pr\big[A(x,y) = 1 \mid (x,y) \in \mathcal{D}\big] - \Pr\big[A(x,y) = 1 \mid (x,y) \notin \mathcal{D}\big], \tag{13}$$

Low $\mathrm{ASR}_{\mathrm{MI}}$ means weak inference; poisoning uses trigger τ to force $y^\star$ in Eq. (14)

$$\mathrm{ASR}_{\mathrm{BD}} = \frac{1}{|\mathcal{D}^\tau|} \sum_{(x,y)\in\mathcal{D}^\tau} \mathbb{I}\big[f_{\mathbf{w}}(x) = y^\star\big], \tag{14}$$

Robust aggregation and anomaly checks enhance security; see Table 6.

Table 6: Summary of attack models and defenses. Inference attacks target privacy and poisoning attacks manipulate model behavior.

Attack	Description	Defenses	Category
Membership inference	Identify if data was in training.	DP, secure aggregation	Inference
Attribute inference	Infer hidden data attributes.	DP noise, regularization	Inference
Label-flip poisoning	Flip labels to mislead training.	Robust aggregation, clipping	Poisoning
Backdoor poisoning	Insert triggers for target outputs.	Trigger checks, robust aggregation	Poisoning
Reconstruction	Recover client data from updates.	Secure aggregation, DP	Reconstruction

3.5 Framework Architecture Overview

Fig. 2 summarizes the privacy-preserving federated AIS analytics pipeline. Raw AIS records (Table 2) are preprocessed into cleaned trajectories and engineered features (Table 3). Samples are grouped into virtual clients using the partitioning schemes in Table 4 and each client trains locally while the server optimizes the global objective with FedAvg (Eq. (10)) under differential privacy and secure aggregation (Table 5). The right panel illustrates the inference and poisoning threats evaluated in Section 3.4 and their defenses (Table 6). In Algorithm 1, the client update $\Delta\mathbf{w}_t^k$ is the same local optimization update defined in Eq. (10) (computed on $\mathcal{D}_k$ and then clipped/noised before aggregation).

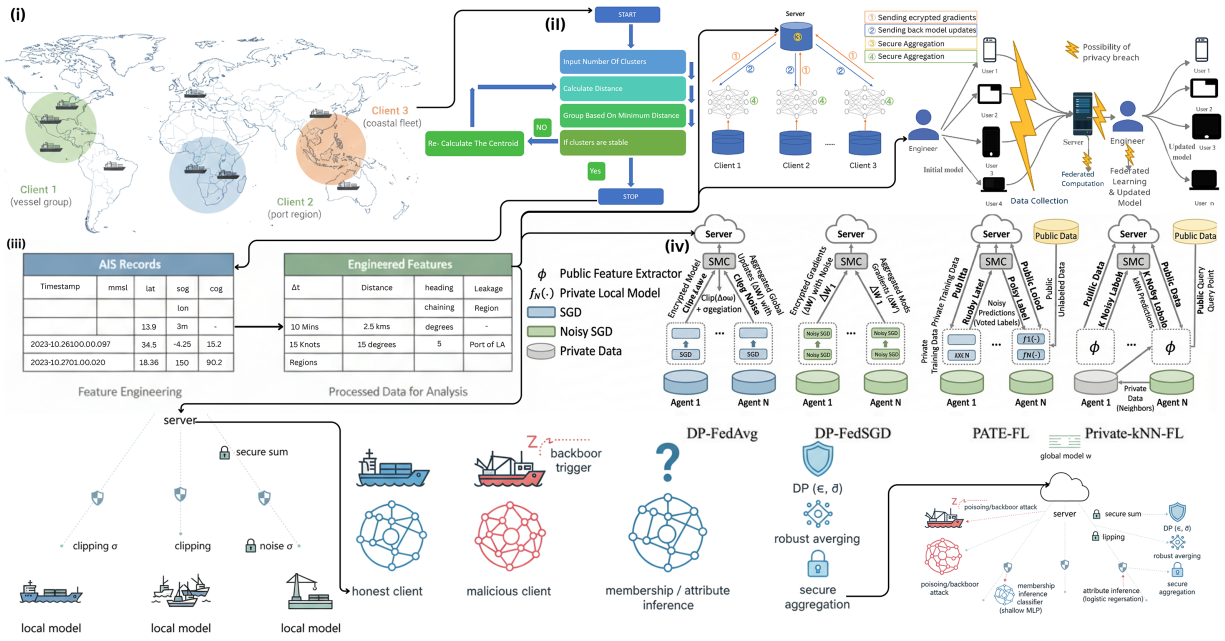

Figure 2: Overview of the privacy-preserving federated AIS analytics framework. From left to right, the figure shows (**i**) vessel- and port-level client construction, (**ii**) AIS feature engineering, (**iii**) federated and differentially private training with secure aggregation, and (iv) attack and defense mechanisms. These four phases correspond to the components in Sections 3.2–3.4.

Algorithm 1: DP-FedAvg with secure aggregation

Initialize global model $\mathbf{w}_0$
for $t = 0$ to $T - 1$ **do**
  Server samples clients $\mathcal{S}_t$ and broadcasts $\mathbf{w}_t$
  for each client $k \in \mathcal{S}_t$ **do**
    Compute local update $\Delta\mathbf{w}_t^k$ on $\mathcal{D}_k$
    Clip and add DP noise to $\Delta\mathbf{w}_t^k$
    Send masked update via secure aggregation
  end for
  Server aggregates masked updates and sets $\mathbf{w}_{t+1}$
end for

3.6 Mapping Privacy Mechanisms to Compliance

To clarify how privacy is achieved without external security services, we map each mechanism to its compliance role shown in Table 7. Privacy is enforced at the algorithmic level through differential privacy, secure aggregation and authenticated participation, ensuring that raw AIS trajectories never leave local clients.

These mappings show that privacy preservation is embedded directly into the learning workflow rather than outsourced to external security services. As a result, the framework maintains confidentiality and regulatory alignment even in zero-trust edge–cloud deployments.

Table 7: Privacy mechanisms aligned with compliance requirements.

Mechanism	What It Protects	Compliance Link
Authentication	Valid client participation	ISO 27001 access control
Secure Aggregation	Hidden client updates	GDPR Art. 32 (secure processing)
Differential Privacy	Non-identifiable trajectories	GDPR Recital 26 (anonymity)
Robust Aggregation	Mitigates poisoning attacks	Operational safety assurance

4 Experiments and Results

This section presents the experimental setup and results.

4.1 Hyperparameter Tuning

Hyperparameters are optimized through mixed search, with search ranges in Table 8.

Table 8: Hyperparameter configuration for federated training, aggregation, and privacy mechanisms.

Hyperparameter	Range/Setting	Value	Hyperparameter	Range/Setting	Value
Learning rate η	$\{10^{-4}, 3\times 10^{-4}, 10^{-3}\}$	10^{-3}	Local batch size B	$\{32, 64, 128, 256\}$	128
Local epochs E	$\{1, 2, 5, 10\}$	2–5	Optimizer	Adam, SGD	Adam
Weight decay λ	$\{0, 10^{-5}, 10^{-4}, 10^{-3}\}$	10^{-4}	Number of rounds T	Fixed	T
Client fraction q	$\{0.1, 0.2, 0.3, 0.5, 1.0\}$	0.2–0.3	Aggregation rule	FedAvg, variants	FedAvg
Client weighting	$\lvert\mathcal{D}_k\rvert$ or uniform	$\lvert\mathcal{D}_k\rvert$	Clipping threshold C	$\{0.1, 0.5, 1.0, 2.0\}$	1.0
Noise multiplier σ	$\{0.5, 0.75, 1.0, 1.5\}$	0.75–1.0	Privacy budget ε	Reported per setting	Reported w/δ
Failure probability δ	fixed	$1/N$	Secure aggregation	Enabled/disabled	Enabled

4.2 Dataset Samples and Client Heterogeneity

Fig. 3 shows substantial label-distribution heterogeneity across clients, with varied Jensen-Shannon divergence values highlighting the strongly non-IID data partitions used in our federated learning experiments.

4.3 Client Label Skew and Secure Aggregation Characteristics

Fig. 4 illustrates pronounced label imbalance across clients, highlighting diverse vessel-type distributions and reinforcing the strongly non-IID characteristics of the federated maritime dataset.

Table 9 shows a 5-row sample illustrating how AIS vessel attributes map to Secure Aggregation payload sizes during federated training. Despite variation in vessel status and characteristics, all clients produce identical 34 kB payloads, reflecting the protocol's fixed-size encrypted updates. This uniformity prevents information leakage through message length and demonstrates communication-level privacy over heterogeneous maritime data.

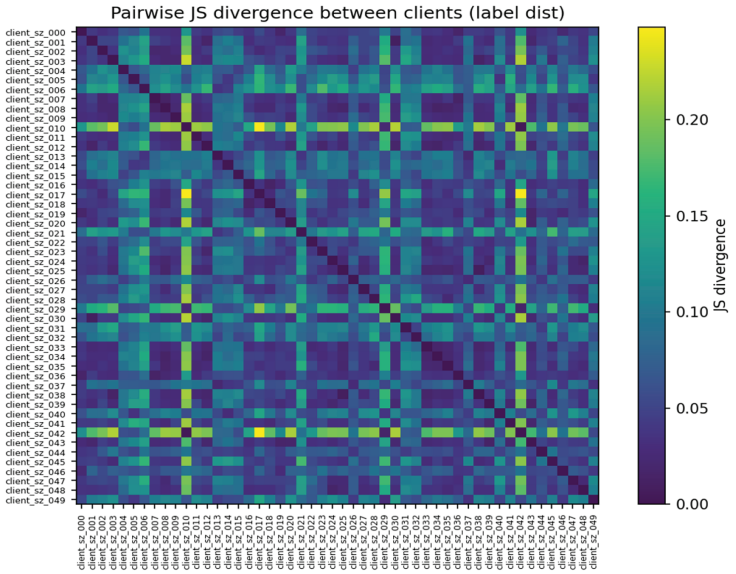

Figure 3: Pairwise JS divergence between clients based on label distributions.

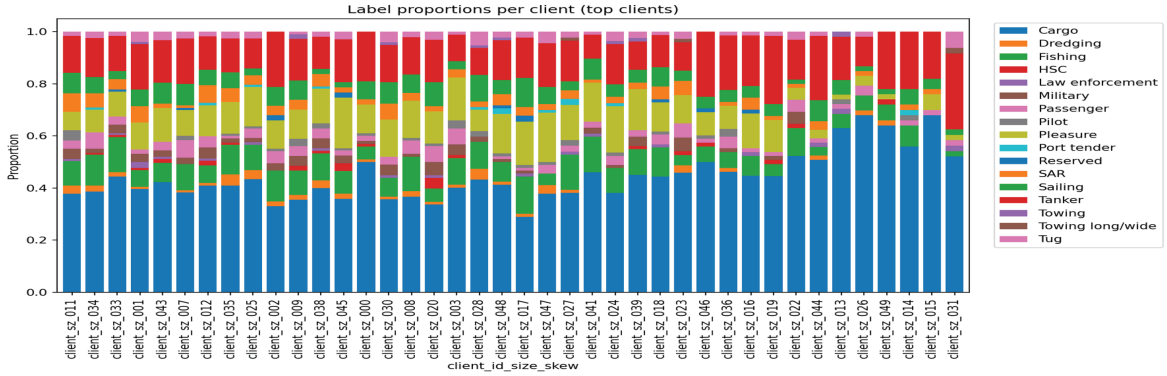

Figure 4: Label distribution proportions for the top clients, illustrating heterogeneity in vessel-type frequencies across the federated network.

Table 9: Random 5-row sample (seed = 42) from `phase3_secureagg_log_full.csv`, showing AIS vessel attributes together with the corresponding Secure Aggregation communication payload size for each record.

Row	Unnamed:0	mmsi	Navigational Status	sog	cog	Heading	Shiptype	Width	Length	Draught	secureagg _bytes_total
2396	123234	257312000	Under way using engine	11.8	134.9	129.0	Cargo	15.0	93.0	5.7	34,000
2323	81133	255924000	Under way using engine	11.3	300.6	300.0	Cargo	13.0	85.0	3.6	34,000
3111	24267	311000280	Under way using engine	8.8	116.0	118.0	Tanker	46.0	276.0	10.0	34,000
1886	116541	244830098	Under way using engine	9.3	344.9	340.0	Cargo	10.0	88.0	4.2	34,000
315	15376	219000158	Moored	0.0	337.3	145.0	Cargo	14.0	99.0	4.5	34,000

4.4 Model Predictions and Utility-Privacy Trade-off

Table 10 shows a sample of AIS records alongside the vessel-type predictions generated by the FedAvg model. The model correctly identifies most Cargo and Tanker vessels, indicating that it captures meaningful motion and dimensional patterns. Occasional errors, often tied to ambiguous behaviour or missing metadata, highlight the challenges of noisy AIS data. The sample demonstrates solid baseline performance while revealing where heterogeneous or sparse client data can limit accuracy.

Table 10: Random 5-row sample (seed = 42) from `phase3_fedavg_preds_full.csv`, showing AIS vessel metadata alongside predicted vessel-type labels produced by the FedAvg model.

Row	Unnamed:0	mmsi	Navigational Status	sog	cog	Heading	Shiptype	Width	Length	Draught	fedavg _pred_label
2396	123234	257312000	Under way using engine	11.8	134.9	129.0	Cargo	15.0	93.0	5.7	Cargo
2323	81133	255924000	Under way using engine	11.3	300.6	300.0	Cargo	13.0	85.0	3.6	Cargo
3111	24267	311000280	Under way using engine	8.8	116.0	118.0	Tanker	46.0	276.0	10.0	Tanker
1886	116541	244830098	Under way using engine	9.3	344.9	340.0	Cargo	10.0	88.0	4.2	Cargo
315	15376	219000158	Moored	0.0	337.3	145.0	Cargo	14.0	99.0	4.5	Fishing

Fig. 5 shows that central and non-private FedAvg models achieve similar macro AUC, while increasing differential privacy noise gradually lowers performance, illustrating the expected utility-privacy trade-off in federated learning.

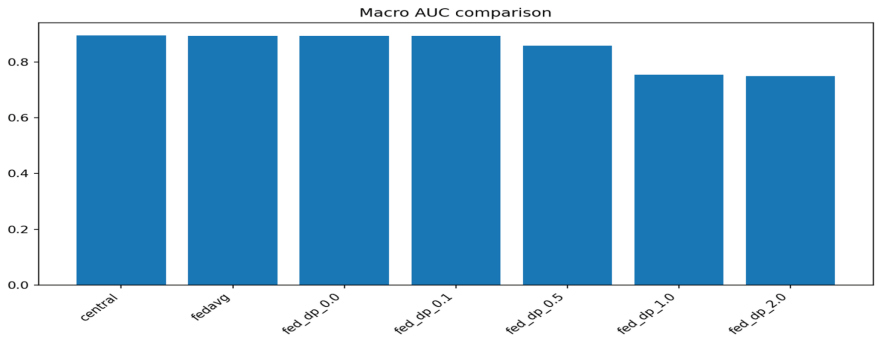

Figure 5: Macro AUC comparison between central, FedAvg and DP-federated configurations.

Fig. 6 shows that Secure Aggregation incurs communication costs similar to FedAvg, while Homomorphic Encryption introduces a substantially higher overhead, illustrating the trade-off between cryptographic strength and system efficiency.

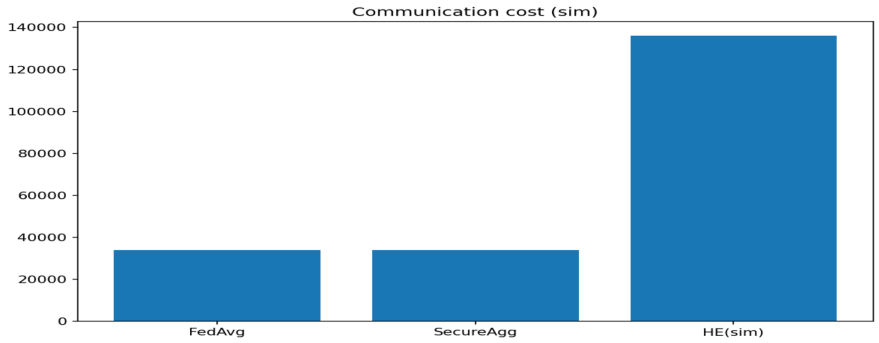

Figure 6: Communication cost comparison across FedAvg, Secure Aggregation and Homomorphic Encryption (HE).

4.5 Communication Overhead and Feature Ablation Inputs

4.6 Ablation Results and Privacy Attack Evaluation

The ablation results in Table 11 show how model performance varies with different differential privacy noise levels and aggregation strategies. When no DP noise is applied ($\sigma = 0$), all aggregation methods achieve similarly strong validation and test AUCs. As noise increases, performance declines gradually, with $\sigma = 1.0$ marking a noticeable reduction across methods. Trimmed-mean aggregation demonstrates slightly greater robustness at moderate noise levels, while extreme perturbation ($\sigma = 5.0$) leads to a substantial collapse in utility. Overall, the tiled results highlight the expected privacy-utility trade-offs and reveal modest differences among the robust aggregation techniques.

Table 11: Ablation results tiled horizontally: 16 entries shown as two side-by-side blocks (each 4 columns).

Block 1				Block 2			
dp_sigma	**agg_method**	**val_auc**	**test_auc**	**dp_sigma**	**agg_method**	**val_auc**	**test_auc**
0.0	avg	0.897436	0.782051	0.1	avg	0.897436	0.782051
0.0	krum	0.897436	0.782051	0.1	krum	0.897436	0.782051
0.0	trimmed	0.897436	0.782051	0.1	trimmed	0.884615	0.769231
0.5	avg	0.884615	0.769231	0.5	krum	0.884615	0.769231
0.5	trimmed	0.871795	0.756410	1.0	avg	0.807692	0.705128
1.0	krum	0.807692	0.705128	1.0	trimmed	0.794872	0.692308
2.0	avg	0.743590	0.641026	2.0	krum	0.743590	0.641026
2.0	trimmed	0.743590	0.641026	5.0	avg	0.256410	NaN

4.7 Poisoning Robustness and Membership Inference Resistance

Fig. 7a shows that increasing the malicious client fraction reduces global accuracy, with median aggregation degrading fastest, while KRUM becomes more resilient under higher adversarial participation. Fig. 7b shows that the membership inference attack performs no better than random guessing, as indicated by an ROC AUC near 0.5, demonstrating minimal privacy leakage.

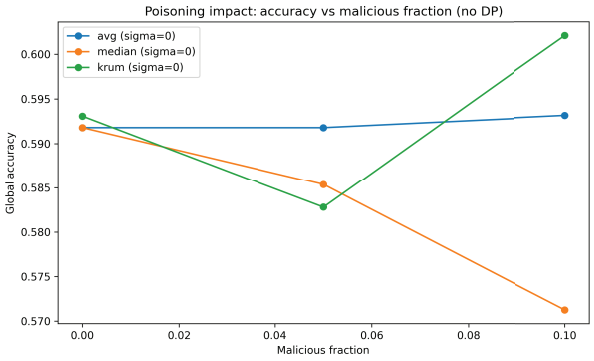

(a) Poisoning impact.

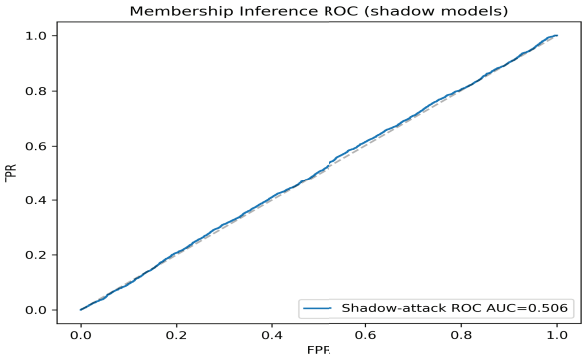

(b) Membership inference.

Figure 7: Poisoning robustness and membership inference resistance.

5 Discussion and Comparison

This section discusses the implications of our findings and compares the proposed framework with existing cloud security and privacy-preserving approaches. Table 12 situates our framework within recent maritime federated learning literature, highlighting differences in tasks, privacy mechanisms and resulting performance.

Table 12: Comparison with recent maritime federated learning studies (post-2023).

Study	Task/Data	FL/Privacy Setup	Main Results
This work	Vessel-type classification on public AIS data (non-IID).	Edge-cloud FedAvg with DP-SGD, secure and robust aggregation.	Macro AUC ≈ 0.90 (val)/0.78 (test); DP-FedAvg (σ=0.5) retains 98% utility; no privacy leakage (ROC AUC ≈ 0.51).
Khodamoradi et al. 2025 [8]	Vessel-density prediction from regional AIS grids.	FedAvg over CNNs, local data privacy only.	Sparse MSE 0.0013–0.0016; no explicit DP or attack defense.
Giannopoulos et al. 2024 [7]	Engine power and fuel-use regression across six ships.	Ship-level FL with ANN; no DP or cryptographic methods.	Accuracy 0.89 vs. 1.00 centralized; MAE 13.3 kW vs. 11.9 kW.

5.1 Utility-Privacy Trade-offs and Non-IID Federated Learning

The empirical results show that our framework maintains strong utility under heterogeneous client data. Exploratory statistics in Tables 2 and 4, together with Figs. 3 and 4, confirm substantial variation and label skew. Despite this, Fig. 5 shows non-private FedAvg remains close to centralized performance, supported by the feature engineering in Section 3.2 and Table 3. Differential privacy degrades performance gradually (Table 11), while secure aggregation incurs minimal overhead (Table 9, Fig. 6). Although validation AUC remains high, the lower test AUC (approximately 0.78) reflects the increased difficulty of fully held-out clients under our non-IID partitioning schemes. In our setting, test clients correspond to vessels or regions with skewed label distributions (Table 5, Fig. 4), so the test performance is intentionally more pessimistic than validation and acts as a proxy for deployment across heterogeneous maritime stakeholders. This gap

therefore indicates realistic generalization challenges in strongly non-IID federated environments rather than overfitting to a single client or centralized split.

5.2 Robustness to Poisoning and Inference Attacks

The robustness results in Section 3.4 show that poisoning can degrade accuracy, with Fig. 7a illustrating median aggregation failing fastest while KRUM remains more resilient. Table 6 highlights the value of combining robust aggregation with DP and secure aggregation. Membership inference scores cluster near 0.35 and Fig. 7b shows an AUC near 0.5, consistent with the clipping and noise in Eqs. (11) and (12). Overall, layered defenses remain essential.

5.3 Practical Implications and Limitations

Beyond quantitative results, the framework provides practical advantages for cloud-enabled maritime systems. Treating vessels and ports as federated clients supports cross-organization learning without exposing raw AIS data, reducing regulatory burdens. The fixed-size encrypted payloads in Table 9 and communication costs in Fig. 6 show feasibility over bandwidth-limited links. Strong FedAvg and DP-FedAvg performance under heterogeneous labels indicates realistic traffic patterns do not hinder training. Remaining limitations include reliance on one AIS dataset, simplified partitions, fixed models and non-adaptive attack assumptions.

6 Conclusion and Future Work

This work presented a privacy-preserving edge-cloud framework for AIS-based maritime analytics using federated learning, differential privacy and secure aggregation. Experiments showed that strong privacy protections can be applied while maintaining competitive predictive performance and communication efficiency. Robust aggregation further mitigated poisoning and inference risks. Overall, the results demonstrate that secure, collaborative maritime intelligence is feasible and that AIS data provides a strong benchmark for privacy-aware learning. Future work includes validating the framework in real maritime deployments, integrating more expressive sequence or graph models and using adaptive privacy or personalized federated optimization to improve utility. Exploring stronger privacy-enhancing technologies, such as MPC or post-quantum cryptography, may further enhance robustness. We also plan to extend the framework to other critical infrastructures.

Acknowledgement: We gratefully acknowledge the AIS dataset.

Funding Statement: This work was supported by the Deanship of Scientific Research, Vice Presidency for Graduate Studies and Scientific Research, King Faisal University, Saudi Arabia Grant No. KFU254769.

Author Contributions: Conceptualization & methodology: Abuzar Khan; Model & experiments: Abid Iqbal; Writing & analysis: Ghassan Husnain; Validation: Fahad Masood; Review: Mohammed Al-Naeem; Supervision: Sajid Iqbal. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: The AIS dataset used in this study is publicly available https://www.kaggle.com/datasets/eminserkanerdonmez/ais-dataset. All resources are available at https://github.com/abuzarkhaaan/A-Secure-and-Differentially-Private-Edge-Cloud.

Ethics Approval: Ethical review not required; AIS data is public dataset.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Ahmad Z, Acarer T, Kim W. Optimization of maritime communication workflow execution with a task-oriented scheduling framework in cloud computing. J Mar Sci Eng. 2023;11(11):2133. doi:10.3390/jmse11112133.
2. Lv T, Tang P, Zhang J. A real-time AIS data cleaning and indicator analysis algorithm based on stream computing. Sci Program. 2023;2023:8345603. doi:10.1155/2023/8345603.
3. Wolsing K, Roepert L, Bauer J, Wehrle K. Anomaly detection in maritime AIS tracks: a review of recent approaches. J Mar Sci Eng. 2022;10(1):112. doi:10.3390/jmse10010112.
4. Chen Y, Zhang G, Liu C, Lu C. Privacy-preserving modeling of trajectory data: secure sharing solutions for trajectory data based on granular computing. Mathematics. 2024;12(23):3681. doi:10.3390/math12233681.
5. Yang Y, Liu Y, Li G, Zhang Z, Liu Y. Harnessing the power of Machine learning for AIS Data-Driven maritime Research: a comprehensive review. Transport Res Part E Logis Transport Rev. 2024;183:103426. doi:10.1016/j.tre.2024.103426.
6. Charpentier V, Slamnik-Kriještorac N, Landi G, Caenepeel M, Vasseur O, Marquez-Barja JM. Paving the way towards safer and more efficient maritime industry with 5G and Beyond edge computing systems. Comput Netw. 2024;250:110499. doi:10.1016/j.comnet.2024.110499.
7. Giannopoulos A, Gkonis PK, Bithas PS, Nomikos N, Ntroulias G, Trakadas P. Federated learning for maritime environments: use cases, experimental results, and open issues. TechRxiv. 2023. doi:10.36227/techrxiv.22133549.
8. Khodamoradi A, Figueiras PA, Grilo A, Lourenço L, Rêga B, Agostinho C, et al. Vessel traffic density prediction: a federated learning approach. ISPRS Int J Geo Inf. 2025;14(9):359. doi:10.3390/ijgi14090359.
9. Zhan S, Huang L, Luo G, Zheng S, Gao Z, Chao HC. A review on federated learning architectures for privacy-preserving AI: lightweight and secure cloud-edge–end collaboration. Electronics. 2025;14(13):2512. doi:10.3390/electronics14132512.
10. Jiao S, Cai L, Meng J, Zhao Y, Cheng K. Efficient DP-FL: efficient differential privacy federated learning based on early stopping mechanism. Comput Syst Sci Eng. 2024;48(1):247–65. doi:10.32604/csse.2023.040194.
11. Guan M, Bao H, Li Z, Pan H, Huang C, Dai HN. SAMFL: secure aggregation mechanism for federated learning with byzantine-robustness by functional encryption. J Syst Archit. 2024;157:103304. doi:10.1016/j.sysarc.2024.103304.
12. Tadi AA, Dayal S, Alhadidi D, Mohammed N. Comparative analysis of membership inference attacks in federated and centralized learning. Information. 2023;14(11):620. doi:10.3390/info14110620.
13. Lim LH, Ong LY, Leow MC. Federated learning for anomaly detection: a systematic review on scalability, adaptability, and benchmarking framework. Future Internet. 2025;17(8):375. doi:10.3390/fi17080375.
14. Li H, Ge L, Tian L. Survey: federated learning data security and privacy-preserving in edge-internet of things. Artif Intell Rev. 2024;57(5):130. doi:10.1007/s10462-024-10774-7.
15. Ribeiro CV, Paes A, de Oliveira D. AIS-based maritime anomaly traffic detection: a review. Expert Syst Appl. 2023;231:120561. doi:10.1016/j.eswa.2023.120561.
16. Xu C, Qu Y, Xiang Y, Gao L. Asynchronous federated learning on heterogeneous devices: a survey. Comput Sci Rev. 2023;50:100595. doi:10.1016/j.cosrev.2023.100595.
17. McMahan B, Moore E, Ramage D, Hampson S, Bay A. Communication-efficient learning of deep networks from decentralized data. In: Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS). Vol. 54. London, UK: PMLR; 2017. p. 1273–82.
18. Dwork C. Differential privacy. In: Bugliesi M, Preneel B, Sassone V, Wegener I, editors. Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006; 2006 Jul 10–14; Venice, Italy. Vol. 4052. Charm, Switzerland: Springer; 2006. p. 1–12. doi:10.1007/11787006.
19. Mohammadi S, Balador A, Sinaei S, Flammini F. Balancing privacy and performance in federated learning: a systematic literature review on methods and metrics. J Parallel Distr Comput. 2024;192:104918. doi:10.1016/j.jpdc.2024.104918.
20. Huang P, Chen Q, Wang D, Wang M, Wu X, Huang X. TripleConvTransformer: a deep learning vessel trajectory prediction method fusing discretized meteorological data. Front Environ Sci. 2022;10:1012547. doi:10.3389/fenvs.2022.1012547.

21. Huang IL, Lee MC, Nieh CY, Huang JC. Ship classification based on AIS data and machine learning methods. Electronics. 2024;13(1):98. doi:10.3390/electronics13010098.
22. Louart M, Szkolnik JJ, Boudraa AO, Le Lann JC, Le Roy F. Detection of AIS messages falsifications and spoofing by checking messages compliance with TDMA protocol. Digit Signal Process. 2023;136:103983. doi:10.1016/j.dsp.2023.103983.
23. Varshitha GS, Rupa C, Divya D, Gadekallu TR, Srivastava G. A survey of authentication protocols for enhancing security in underwater communication systems. In: 2025 IEEE 34th Wireless and Optical Communications Conference (WOCC). Piscataway, NJ, USA: IEEE; 2025. p. 1–6.
24. Rupa C, Varshitha GS, Divya D, Gadekallu TTR, Srivastava G. A novel and robust authentication protocol for secure underwater communication systems. IEEE Inter Things J. 2025;12(22):47519–31. doi:10.1109/jiot.2025.3601984.
25. Barona López LI, Borja Saltos T. Heterogeneity challenges of federated learning for future wireless communication networks. J Sens Actuator Netw. 2025;14(2):37. doi:10.3390/jsan14020037.
26. Tayyeh HK, AL-Jumaili ASA. Balancing privacy and performance: a differential privacy approach in federated learning. Computers. 2024;13(11):277. doi:10.3390/computers13110277.
27. Zhang X, Luo Y, Li T. A review of research on secure aggregation for federated learning. Future Inter. 2025;17(7):308. doi:10.3390/fi17070308.
28. Erdonmez ES. AIS dataset: transition of ships at kattegat strait; 2020. Kaggle Dataset. [cited 2026 Jan 10]. Available from: https://www.kaggle.com/datasets/eminserkanerdonmez/ais-dataset.