



MH-GIN: Multi-scale Heterogeneous Graph-based Imputation Network for AIS Data

Hengyu Liu
Aalborg University, Denmark
heli@cs.aau.dk

Tianyi Li
Aalborg University, Denmark
tianyi@cs.aau.dk

Yuqiang He
Guangxi University, China
heyq1314@gmail.com

Kristian Torp
Aalborg University, Denmark
torp@cs.aau.dk

Yushuai Li
Aalborg University, Denmark
yusli@cs.aau.dk

Christian S. Jensen
Aalborg University, Denmark
csj@cs.aau.dk

ABSTRACT

Location-tracking data from the Automatic Identification System, much of which is publicly available, plays a key role in a range of maritime safety and monitoring applications. However, the data suffers from missing values that hamper downstream applications. Imputing the missing values is challenging because the values of different heterogeneous attributes are updated at diverse rates, resulting in the occurrence of multi-scale dependencies among attributes. Existing imputation methods that assume similar update rates across attributes are unable to capture and exploit such dependencies, limiting their imputation accuracy. We propose MH-GIN, a Multi-scale Heterogeneous Graph-based Imputation Network that aims to improve imputation accuracy by capturing multi-scale dependencies. Specifically, MH-GIN first extracts multi-scale temporal features for each attribute while preserving their intrinsic heterogeneous characteristics. Then, it constructs a multi-scale heterogeneous graph to explicitly model dependencies between heterogeneous attributes to enable more accurate imputation of missing values through graph propagation. Experimental results on two real-world datasets find that MH-GIN is capable of an average 57% reduction in imputation errors compared to state-of-the-art methods, while maintaining computational efficiency.

PVLDB Reference Format:

Hengyu Liu, Tianyi Li, Yuqiang He, Kristian Torp, Yushuai Li, and Christian S. Jensen. MH-GIN: Multi-scale Heterogeneous Graph-based Imputation Network for AIS Data. PVLDB, 19(2): 170 - 182, 2025.
doi:10.14778/3773749.3773756

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/hyLiu1994/MH-GIN>.

1 INTRODUCTION

The Automatic Identification System (AIS) is an automated tracking system that enhances navigational safety by enabling vessels to share their position, identification, and other essential information with nearby ships and coastal authorities [3]. AIS data is crucial for

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment. Tianyi Li is the corresponding author. Proceedings of the VLDB Endowment, Vol. 19, No. 2 ISSN 2150-8097.
doi:10.14778/3773749.3773756

Index	Observed Attribute		Missing Attribute		Update rates						
	Longitude	Latitude	COG	SOG	Navi.	Draught	Length	Width	Vessel Type	MMSI	
1	57.694620°	10.619458°	186	?	Underway using engine	15	?	?	Tanker	20187312	
2	?	?	185	15	Underway using engine	15	?	?	Tanker	20187312	
3	57.454070°	10.563542°	0	?	Anchor	15	?	?	Tanker	20187312	
4	57.454071°	10.5635421°	0	0	Anchor	15	?	?	Tanker	20187312	
5	57.454069°	10.563542°	?	0	Anchor	15	?	?	Tanker	20187312	
6	?	?	175	3	?	18	?	?	Tanker	20187312	
7	57.449921°	10.566126°	175	10	?	18	?	?	Tanker	20187312	
8	57.449921°	10.566113°	?	10	?	18	?	?	Tanker	20187312	

Figure 1: An example of AIS records from a vessel.

maritime stakeholders to operate applications spanning different domains, including vessel tracking, safety monitoring, logistics optimization, and trade forecasting. However, AIS data often suffer from quality issues, among which missing values represent a significant challenge. Missing values may be caused by a variety of various factors, such as equipment malfunctions, vessels switching off their transponders, or vessels being out of range of terrestrial receivers [25]. Notably, statistics of data from the Danish Maritime Authority [10] reveal that missing rates in attributes range from 8% to 83%, which underscores the substantial need for data imputation in real-world scenarios. When AIS fails, downstream applications (such as vessel tracking [51], safety monitoring [46], and navigation planning [38]) often rely on traditional tools (such as radar and VHF radio [42, 43]) or alternative systems (such as Long-Range Identification and Tracking (LRIT) and Vessel Monitoring Systems (VMS) [15]) typically at the cost of accuracy and efficiency.

Example 1. Consider a vessel's AIS records during a port visit as illustrated in Figure 1. The vessel follows a typical maritime pattern: entering the port (records 1–2), berthing at the terminal (records 3–5), and departing from the port (records 6–8). During this process, the vessel's Longitude, Latitude, COG, and SOG are updated at high frequency and exhibit random missing patterns. The vessel's Navigation Status and Draught are updated at low frequency and show block missing patterns. The vessel's Length, Width, and Vessel Type remain constant throughout, with Length and Width being entirely missing.

As shown in Example 1, AIS data exhibit three characteristics [39]: 1) There are **heterogeneous attributes** in AIS data, i.e., spatio-temporal attributes (Longitude, Latitude), discrete attributes

Table 1: Summary of existing imputation methods

Methodology Class	Heterogeneous Attributes	Diverse Update Rates	Multi-scale Dependencies
Heuristic & Statistical Imputation	Yes	No	No
Multi-variable Time Series Imputation	No (Assumes Homogeneity)	No (Assumes Uniform Update Rates)	No (Single-scale Dependency)
Spatio-temporal Imputation	Partial (Focus on Spatial and Continuous Attributes)	Partial (Assumes Uniform Rates for Non-spatial)	Partial (Single-scale Dependency for Non-spatial)
Trajectory Imputation	No (Spatio-temporal Attributes Only)	No (Assumes Similar Update Rates)	No (Single-scale Dependency)

(Navigation Status), and continuous attributes (Draught). 2) Different attributes exhibit **diverse update rates**, leading to diverse missing patterns. As the update rate decreases, the missing pattern transitions from random to block to entirely missing. 3) There are **multi-scale dependencies** between attributes. Attributes at high time scales constrain the high-level features of attributes at low time scales without affecting their low-level features. Conversely, only the high-level features of attributes at low time scales influence attributes at higher time scales. For instance, as shown in Figure 1, Navigation Status (Anchored) constrains a vessel’s movement to a limited region around a fixed point (high-level feature: movement pattern of Longitude and Latitude) without specifying precise coordinates. Similarly, Vessel Type (Tanker) constrains the value range of Draught (high-level feature of Draught) without specifying the exact value. Conversely, the value range of Draught also constrains the value range of the vessel’s Length and Width.

However, existing imputation methods cannot effectively handle these characteristics. As shown in Table 1, existing imputation methods can be categorized into four categories, each with distinct limitations: 1) Traditional and statistical methods [1, 2, 5, 17, 27, 36, 49] fall short due to their assumptions of linear relationships and their inability to handle diverse update rates and model multi-scale dependencies. 2) Multi-variable time series imputation methods [7, 8, 11, 40] assume uniform update rates and treat attributes homogeneously. 3) Spatio-temporal imputation methods [24, 26, 31] assume static spatial relationships and uniform update rates for non-spatial attributes. 4) Trajectory imputation methods [30, 35, 55] focus exclusively on reconstructing vessel positions and rarely consider other heterogeneous attributes critical for comprehensive AIS analysis. Effectively handling the characteristics of AIS data is challenging and presents significant obstacles.

C1: How to effectively represent and simultaneously impute heterogeneous attributes? AIS data consist of four attribute types: spatio-temporal, cyclical, continuous, and discrete, and each has distinct intrinsic characteristics. The heterogeneous attributes require unified yet type-specific representation and imputation strategies, capable of preserving their individual properties. Such strategies are an essential precondition for capturing multi-scale dependencies and accurately imputing missing data. Existing imputation methods [8, 26, 31, 40, 55] exhibit significant limitations in addressing this heterogeneity. They either focus exclusively on the spatio-temporal data or treat all attributes homogeneously, neglecting the unique properties of different attribute types.

C2: How to effectively extract multi-scale temporal features from each attribute? As discussed earlier, AIS attributes operate

at different time scales and influence each other indirectly. To capture the direct interactions among attributes at a particular scale accurately, it is essential to model temporal features individually at their own and higher scales. However, existing imputation methods [6, 8, 11, 40] typically assume similar update rates across all attributes, thus extracting temporal features indiscriminately at all scales into a single, uninterpretable feature vector.

C3: How to effectively model multi-scale dependencies between attributes? Multi-scale dependencies are complex, involving both dependencies between attributes at the same time scale and dependencies between attributes across different time scales. Effectively modeling and leveraging these dependencies for accurate imputation is challenging. Current methods [4, 7, 11, 53] do not differentiate direct and indirect attribute interactions at varying scales due to their similar update rate assumption, resulting in oversimplified attribute relationships and reduced imputation performance.

To address these challenges, we propose MH-GIN, a Multi-scale Heterogeneous Graph-based Imputation Network for AIS data.

Addressing C1: MH-GIN utilize type-specific encoding and decoding strategies for four distinct attribute types: spatio-temporal, cyclical, continuous, and discrete attributes, while preserving their intrinsic characteristics, such as spatial constraints, smooth wrap-around at boundaries.

Addressing C2: MH-GIN adopts a hierarchical temporal feature extractor based on Deep Echo State Networks, enabling the capture of hierarchical temporal features for each attribute through progressive abstraction and leaky integration.

Addressing C3: MH-GIN constructs a multi-scale heterogeneous graph to explicitly model two types of multi-scale dependencies. Specifically, the graph contains two kinds of subgraphs: time-scale subgraphs and attribute subgraphs. The former model dependencies between features of different attributes at the same scale, whereas the latter model dependencies across different scales within each attribute. Subsequently, MH-GIN imputes missing features through a two-stage graph propagation process, first aligning features within time-scale subgraphs, and then impute missing features through propagation within attribute subgraphs.

Our main contributions are summarized as follows:

- We propose MH-GIN, a Multi-scale Heterogeneous Graph-based Imputation Network for AIS data, which captures multi-scale dependencies among heterogeneous attributes to enhance imputation performance.
- We design complementary modules, including type-specific encoders and decoders and a hierarchical temporal feature extractor, to collaboratively extract hierarchical temporal features for each attribute type while preserving intrinsic characteristics of each attribute type.
- We introduce a multi-scale dependency mining module, which constructs a multi-scale heterogeneous graph to model dependencies between attributes explicitly and performs imputation via two-stage graph propagation.
- We evaluate MH-GIN extensively on two real-world AIS datasets, finding that MH-GIN is able to outperform state-of-the-art approaches, achieving an average 57% reduction in imputation error while maintaining computational efficiency.

The remainder of the paper is structured as follows: Section 2 reviews related work. Section 3 presents preliminaries and formulates the problem. Section 4 details the proposed MH-GIN. Section 5 reports experimental results. Section 6 concludes the paper and outlines directions for future work.

2 RELATED WORK

Existing imputation methods can be categorized into four categories, as shown in Table 1, each with distinct limitations:

Traditional and Statistical Imputation Methods. Traditional approaches for handling missing data include deletion methods, neighbor-based methods, constraint-based methods, and statistical imputation. Deletion methods, such as listwise deletion [49] and pairwise deletion [27], directly remove data instances containing missing values, often resulting in significant information loss. Neighbor-based methods like KNN [5] impute missing values using local patterns derived from similar data points. Constraint-based methods [36] leverage domain-specific rules, performing well when data adheres to particular constraints or patterns. Common statistical approaches include mean/median imputation [1], Last Observation Carried Forward (LOCF)[2], and linear interpolation[17]. Despite being computationally efficient and straightforward, these traditional methods are inadequate for AIS data due to their implicit assumptions of linear relationships. AIS data inherently exhibits complex, non-linear patterns. Additionally, deletion-based approaches are particularly problematic for AIS, given its often extended periods of missingness. Thus, traditional methods struggle to represent the inherent multi-scale dependencies and heterogeneous attributes within AIS datasets.

Multi-variable Time Series Imputation. Deep learning methods for multi-variable time series imputation have attracted considerable attention in recent years. Early RNN-based methods, such as GRU-D [7], leverage gating mechanisms and decay factors to handle irregularly sampled data. More recently, transformer-based architectures, including SAITS [11] and DeepMVI [4], utilize self-attention mechanisms to effectively capture long-range dependencies. Diffusion-based methods, such as CSIDI [40], employ conditional score-based diffusion models to produce high-quality imputations through iterative data denoising. Additionally, graph-based methods like GRIN [8] represent each variable as a graph node, enabling message passing to capture cross-dimensional correlations. However, existing general-purpose imputation methods typically assume uniform update rates across all variables and treat attributes homogeneously. These assumptions conflict with the inherent characteristics of AIS data, which contains heterogeneous attributes operating at multiple time scales.

Spatio-temporal Imputation. Spatio-temporal imputation extends conventional time series methods by explicitly modeling spatial dependencies (e.g., sensor grids or geographic adjacency). Recent approaches incorporate graph neural networks (GNNs) [50] or attention mechanisms to jointly learn spatio-temporal patterns. Representative methods include PriSTI [24], which enhance diffusion-based models by capturing spatio-temporal dependencies and geographic relationships. ImputeFormer [31] utilizes a low-rank Transformer architecture to achieve a balance between inductive bias

and model expressivity. SPIN [26] proposes a spatio-temporal propagation framework with sparse attention, specifically designed to handle highly sparse observations by conditioning reconstruction exclusively on available data, thus preventing error propagation common in autoregressive GNNs. Although effective in sensor networks with relatively fixed spatial structures, these methods generally assume stationary spatial relationships or static graph topologies. However, AIS data involves vessels continuously moving across large maritime regions, leading to dynamic and evolving spatial relationships rather than static sensor locations. Therefore, existing spatio-temporal imputation methods face challenges in effectively modeling these dynamic spatial relationships. Furthermore, similar to multi-variable time series imputation methods, they typically assume uniform update rates for non-spatial attributes, rendering them ineffective in capturing diverse update frequencies inherent to AIS attributes.

Trajectory Imputation. Trajectory imputation methods reconstruct missing segments in the trajectories of moving objects (e.g., vehicles, pedestrians, and ships), typically leveraging motion continuity and spatio-temporal context. Transformer-based approaches, such as TrajBERT [35], embed geographic coordinates as tokens and reconstruct missing positions using bidirectional attention mechanisms. KAMEL [28] maps trajectory imputation to the missing word problem in natural language processing, adapting BERT with spatial-awareness and multi-point imputation capabilities for scalable trajectory reconstruction. Recurrent and graph-based sequence models [22, 52, 53, 56] exploit temporal dependencies and spatial constraints, capturing complex motion patterns and environmental context to impute missing trajectory segments. Meanwhile, generative models like GANs [34] and VAEs [33] produce plausible trajectories consistent with known movement behaviors. Specifically tailored maritime approaches include Multi-task AIS [30], which integrates recurrent neural networks with latent-variable modeling to simultaneously address trajectory reconstruction, anomaly detection, and vessel-type identification. PG-DPM [55] introduces a physics-guided diffusion probabilistic model specifically designed for long-term vessel trajectory imputation by incorporating maritime domain knowledge into the generative process. Complementary studies on trajectory compression and clustering [18–20, 37] share similar goals of preserving motion patterns and data utility. However, these methods generally focus exclusively on reconstructing vessel positions and rarely consider other heterogeneous attributes, which are critical for comprehensive AIS data analysis.

3 PRELIMINARIES

3.1 Data and Notation

Definition 1 (AIS Record). An AIS Record X with N attributes can be characterized from two perspectives (see Table 2): (i) **attribute type**: including spatio-temporal set $X_s = \{\lambda, \phi, \tau\}$, cyclical set $X_c = \{\psi, \theta\}$, continuous set $X_n = \{s, d, \ell, \beta\}$, and discrete set $X_d = \{\eta, \chi, \kappa\}$; (ii) **time scale (update rate)**: attributes are grouped according to decreasing update rates into five time scales: $X^1 = \{\lambda, \phi, \tau\}$, $X^2 = \{\psi, \theta, s\}$, $X^3 = \{\eta\}$, $X^4 = \{\chi, d\}$, and $X^5 = \{\ell, \beta, \kappa\}$.

An attribute can be denoted by its attribute type and time scale. For instance, Speed Over Ground (SOG) is a continuous attribute

Table 2: Summary of AIS attributes categorized by type and time scale.

Attribute (Symbol)	Attribute Type				Time Scale				
	Spatio-temporal (X_s)	Cyclical (X_c)	Continuous (X_n)	Discrete (X_d)	$S1 (X^1)$	$S2 (X^2)$	$S3 (X^3)$	$S4 (X^4)$	$S5 (X^5)$
Longitude (λ)	✓				✓				
Latitude (ϕ)	✓				✓				
Timestamp (τ)	✓				✓				
True heading angle (ψ)		✓				✓			
Course Over Ground (θ)		✓				✓			
Speed Over Ground (s)			✓			✓			
Navigation status (η)				✓			✓		
Hazardous cargo type (χ)				✓				✓	
draught (d)				✓				✓	
Length (ℓ)				✓					✓
Width (β)				✓					✓
Vessel type (κ)					✓				✓

[†]Si represents time scale i , where $i \in \{1, 2, 3, 4, 5\}$

[‡] Attribute types and time scales are organized according to Definitions 3 and 4.

x_n^2 at time scale 2, denoted by s . It indicates the vessel's speed, and simultaneously belongs to X^2 , X_n , and X_n^2 .

Definition 2 (Vessel-specific AIS Record Sequence). A vessel-specific AIS record sequence is a heterogeneous multi-variable time series denoted by $\mathbf{X} = \langle X_1, X_2, \dots, X_T \rangle \in \mathbb{R}^{T \times N}$. Each record X_t at time step t can be decomposed into two complementary views: by time scale (update rate), $X_t = \{X_t^1, X_t^2, X_t^3, X_t^4, X_t^5\}$; by attribute type, $X_t = \{X_{t,s}, X_{t,c}, X_{t,n}, X_{t,d}\}$.

For example, SOG s_t at time step t is a continuous attribute $x_{n,t}^2$ at time scale 2, belonging to $X_{t,n}^2$, $X_{t,n}$, and X_t^2 .

AIS data attributes can generally be categorized into three types based on their intrinsic characteristics: spatio-temporal, continuous, and discrete. To achieve a finer granularity and better applicability, we introduce a fourth category by extracting COG θ and THA ψ from the continuous attributes into cyclical attributes due to their periodic nature. The final categorization is as follows:

Definition 3 (Intrinsic Characteristics of AIS Attributes). AIS attributes can be categorized into four types based on their intrinsic characteristics, as summarized in Table 2:

- **Spatio-temporal** attributes exhibit spatial constraints (e.g., latitude-dependent distortion and spherical continuity) and nested periodicities (e.g., daily, weekly cycles).
- **Cyclical** attributes recur periodically and exhibit smooth wrap-around at boundaries.
- **Continuous** attributes exhibit wide variations in scale, requiring appropriate normalization.
- **Discrete** attributes belong to finite categorical sets without inherent numerical ordering.

Spatio-temporal attributes exhibit spatial and temporal characteristics. For instance, spatial scaling varies with latitude: at the equator ($\phi = 0^\circ$), a longitudinal change of 1° corresponds to approximately 111 km, while at a latitude of $\phi = 79^\circ$ North, the same 1° shift corresponds to roughly 20 km. Moreover, longitude values of $\lambda = 180^\circ$ and $\lambda = -180^\circ$ represent the same geographic meridian, due to spherical continuity. Temporally, timestamps demonstrate nested periodicities – both daily (24×3600 seconds) and weekly (168×3600 seconds) cycles return to the starting point ($\tau = 0$).

Cyclical attributes, such as heading angle (ψ), measured from 0° to 359° , demonstrates circular continuity, where angles exceeding 359° wrap around to 0° without causing discontinuity.

Continuous attributes, such as draught (d) and speed over ground (s), differ substantially in their numerical ranges. Draught typically ranges from 0 to 20 meters, while speed ranges from 0 to 30 knots or higher. Due to the distinct scales, a unit change in draught and a unit change in speed are not directly comparable, thus requiring normalization prior to unified analyses.

Discrete attributes represent categorical values without inherent numerical order. For example, the navigation status attribute (η) includes distinct categories such as 0 (under way), 1 (anchored), and 2 (not under command). These values function purely as categorical labels rather than as numerical measurements.

3.2 Multi-scale Heterogeneous Graph

AIS data attributes exhibit distinct update patterns that naturally lead to five time scales. According to the update mechanism [25], attributes can initially be categorized into three groups: (i) Autonomous sensor updates for the first six attributes in Table 2; (ii) Periodic crew updates for the seventh and eighth attributes; (iii) One-time vessel registration updates for the last four attributes.

To capture finer temporal granularities, we further subdivide these groups: autonomous sensor attributes are split based on whether they require measurement intervals, while crew-updated attributes are divided based on whether they only change between voyages. This results in five distinct time scales that reflect the natural update frequencies of AIS attributes.

Definition 4 (Multiple Time Scales). We categorize AIS data attributes into five distinct time scales based on their update rates, as summarized in Table 2. Specifically:

- **Time scale 1:** Attributes $x^1 \in X^1$ have instantaneous values without requiring measurement intervals; thus, their theoretical update interval lower bound is zero.
- **Time scale 2:** Attributes $x^2 \in X^2$ change rapidly and require short measurement intervals, resulting in a theoretical lower bound slightly above zero.
- **Time scale 3:** Attributes $x^3 \in X^3$ change discretely at a moderate frequency during a voyage; hence, their update intervals have a theoretical lower bound above that of time scale 2.
- **Time scale 4:** Attributes $x^4 \in X^4$ remain constant within individual voyages but may change between voyages, setting their theoretical lower bound to the duration of the shortest voyage.

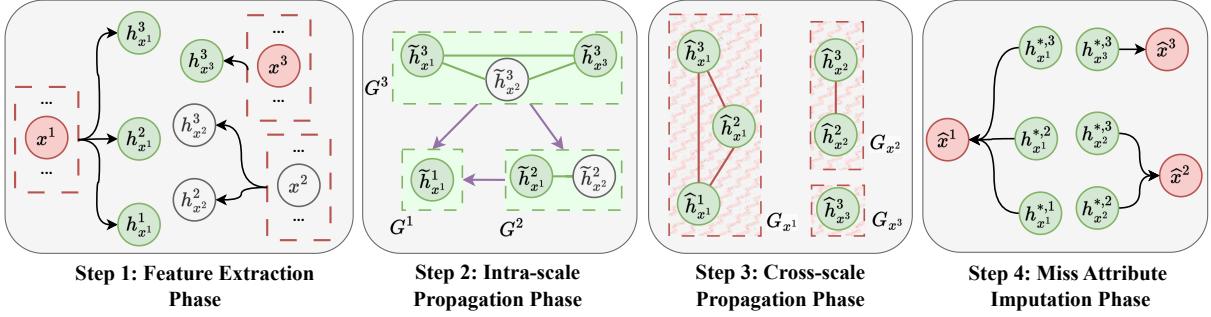


Figure 2: An example of missing attribute imputation based on a multi-scale heterogeneous graph.

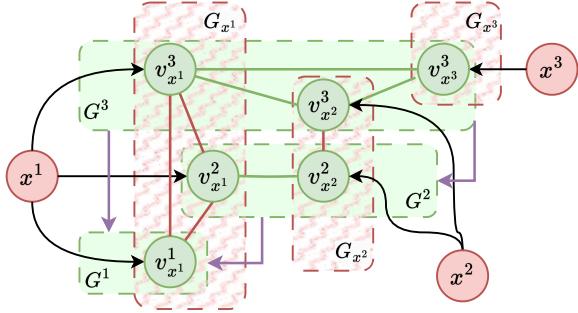


Figure 3: An example of multi-scale heterogeneous graph.

- **Time scale 5:** Attributes $x^5 \in X^5$ change rarely during a vessel’s operational lifespan, resulting in a very large theoretical lower bound for their update interval.

For example, SOG is classified as time scale 2, as it requires short intervals to accurately measure its rapidly changing value.

Definition 5 (Multi-scale Temporal Features). AIS data attributes generate hierarchical features across different time scales. Specifically, for an attribute $x^j \in X^j$ at its original time scale j , features $v_{x^j}^k \in \mathcal{V}^k$ can be derived at time scales k , where $k \geq j$. Thus, each attribute x^j generates $5 + 1 - j$ multi-scale features that capture its temporal features at different time scales.

Example 2. Consider three attributes x^1 , x^2 , and x^3 at time scales 1, 2, and 3, respectively (red nodes in Figure 3). Each attribute generates temporal features at its own and higher time scales: x^1 generates features $(v_{x^1}^1, v_{x^1}^2, v_{x^1}^3)$, x^2 generates features $(v_{x^2}^2, v_{x^2}^3)$, and x^3 generates the feature $v_{x^3}^3$ (green nodes in Figure 3).

Definition 6 (Multi-scale Heterogeneous Graph). A multi-scale heterogeneous graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ represents temporal features derived from all attributes across multiple time scales. Specifically, the node set \mathcal{V} comprises the temporal features, while the edge set \mathcal{E} captures dependencies between them. The graph structure integrates two types of relationships: (i) **Time scale subgraphs** $\{\mathcal{G}^k\}_{k=1}^5$: each subgraph $\mathcal{G}^k = (\mathcal{V}^k, \mathcal{E}^k)$ is fully connected, consisting of nodes \mathcal{V}^k denoting features at the same time scale k as edges \mathcal{E}^k , capturing intra-scale dependencies; edges \mathcal{E}^k are informed by higher-scale subgraphs $\mathcal{G}^{k+1}, \dots, \mathcal{G}^5$, ensuring that relationships at lower time scales are contextually modulated by high-scale features; (ii) **Attribute subgraphs** $\{\mathcal{G}_x\}_{x \in X}$: each subgraph $\mathcal{G}_x = (\mathcal{V}_x, \mathcal{E}_x)$ is fully connected, containing nodes denoting features of a specific attribute x across multiple time scales. These attribute subgraphs are mutually

independent, and edges \mathcal{E}_x capture cross-scale dependencies within the same attribute.

Example 3. Continuing Example 2, feature $v_{x^1}^2$ establishes intra-scale connections with $v_{x^2}^2$ through subgraph \mathcal{G}^2 (shown in the green box labeled \mathcal{G}^2), and the edge between $v_{x^1}^2$ and $v_{x^2}^2$ is modulated by higher-scale subgraphs \mathcal{G}^3 (shown in purple lines from \mathcal{G}^3 to \mathcal{G}^2). Simultaneously, $v_{x^1}^2$ connects across scales within its attribute subgraph $\mathcal{G}x^1$ (shown as the red box labeled $\mathcal{G}x^1$), linking features $v_{x^1}^1$, $v_{x^1}^2$, and $v_{x^1}^3$. However, $v_{x^1}^2$ does not connect to $v_{x^2}^3$ due to differences in both attribute type and time scale.

3.3 Problem Formulation

Definition 7 (Observation Mask). Observation availability is represented by a binary mask $M \in \{0, 1\}^{T \times N}$, where $m_{xt} = 1$ indicates that attribute x is observed at time step t , and $m_{xt} = 0$ indicates that the attribute is missing. For training and evaluation purposes, we select imputation targets $\tilde{X} \in \mathbb{R}^{T \times N}$ manually from the observed data and mark these selected positions using mask $\tilde{M} \in \mathbb{R}^{T \times N}$.

The masking pattern of \tilde{M} varies across the attributes at different time scales in a record sequence: attributes X^5 at time scale 5 generally exhibit either complete availability or are missing extensively; attributes X^4 at time scale 4 present inter-voyage missing patterns, changing values mainly between voyages; attributes X^3 at time scale 3 experience missing data related specifically to voyage phases; attributes X^2 and X^1 at time scales 1 and 2 frequently suffer from irregular and intermittent missing values due to signal interference or communication issues. The details of the masking pattern are discussed in Section 5.1.4.

Definition 8 (Multi-scale Heterogeneous Graph-based Missing Attribute Imputation). The imputation process is formulated as a function \mathcal{F} that maps a record sequence $X * M$ with missing attributes to a complete record sequence \tilde{X} :

$$\tilde{X} = \mathcal{F}(X * M), \quad (1)$$

where \tilde{X} is the imputed record sequence. Specifically, $\mathcal{F}(\cdot)$ can be divided into four phases:

- **Feature Extraction Phase** ($H \leftarrow \mathcal{F}_1(X * M)$): Extract temporal representations H of all attributes at distinct scales from record sequence $X * M$ with missing attributes.
- **Intra-scale Propagation Phase** ($\tilde{H} \leftarrow \mathcal{F}_2(H)$): Propagate representations \tilde{H} within time scale subgraphs $\{\mathcal{G}^k\}_{k=1}^5$ to ensure temporal consistency within each scale. In this process, the edges

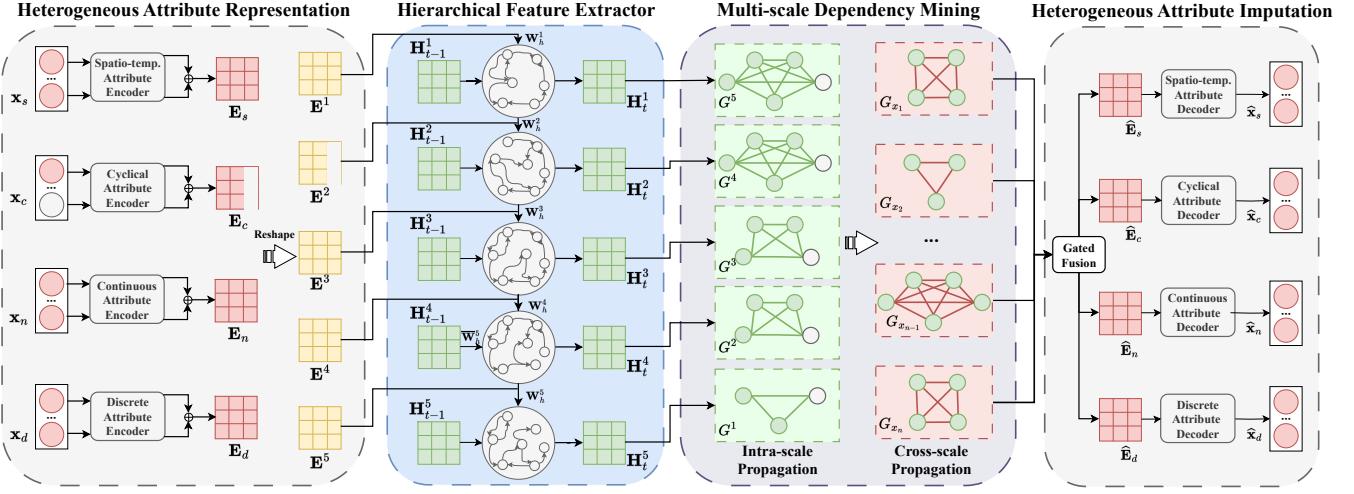


Figure 4: Overview of the multi-scale heterogeneous graph imputation network for AIS data.

$\{\mathcal{E}^k\}_{k=1}^5$ within each time scale subgraph is modulated by the higher-scale subgraphs, ensuring that the relationships at lower time scales are contextually modulated by high-scale features.

- **Cross-scale Propagation Phase** ($\hat{\mathbf{H}} \leftarrow \mathcal{F}_3(\tilde{\mathbf{H}})$): Propagate representations $\hat{\mathbf{H}}$ across different time scales within attribute subgraphs $\{\mathcal{G}_x\}_{x \in \mathcal{X}}$ to impute missing features.
- **Missing Attribute Imputation Phase** ($\tilde{\mathbf{X}} \leftarrow \mathcal{F}_4(\mathbf{H}, \tilde{\mathbf{H}}, \hat{\mathbf{H}})$): Recover missing attributes $\tilde{\mathbf{X}}$ based on \mathbf{H}^* that combines the original representations \mathbf{H} and propagated representations $\tilde{\mathbf{H}}$ and $\hat{\mathbf{H}}$ from intra-scale and cross-scale propagation.

Example 4. Continuing Example 2, as seen in Figure 2, the imputation for the missing attribute x^2 begins with multi-scale representation extraction. Attributes at different scales produce representations: attribute x^1 generates representations $h_{x^1}^1, h_{x^1}^2, h_{x^1}^3$ (scales 1–3), attribute x^3 produces $h_{x^3}^3$ (scale 3), while attribute x^2 yields incomplete representations $h_{x^2}^2, h_{x^2}^3$ (scales 2–3). Temporal consistency is maintained through intra-scale propagation on subgraphs $\mathcal{G}^1, \mathcal{G}^2$, and \mathcal{G}^3 (see the green box). The intra-scale propagation employs dynamic relationship modeling where edge weights in lower-scale subgraphs (e.g., \mathcal{G}^2) are dynamically modulated by features from higher-scale subgraphs (e.g., \mathcal{G}^3), enabling contextual adaptation of attribute relationships across different temporal scales. This results in refined representations $\{\tilde{h}_{x^1}^k\}_{k=1}^3, \{\tilde{h}_{x^3}^k\}_{k=1}^3$, and partially recovered representations $\{\tilde{h}_{x^2}^k\}_{k=2}^3$. Subsequent cross-scale propagation within attribute subgraphs $\mathcal{G}_{x^1}, \mathcal{G}_{x^2}$, and \mathcal{G}_{x^3} (see the red box) synthesizes missing representations $\tilde{h}_{x^2}^2$ and $\tilde{h}_{x^2}^3$ through hierarchical representation fusion. The final imputation integrates original, intra-scale, and cross-scale representations as $\tilde{h}_{x^2}^{*,k} = [h_{x^2}^k, \tilde{h}_{x^2}^k, \hat{h}_{x^2}^k]$ for $k \in \{1, 2\}$, ultimately reconstructing the attribute \tilde{x}^2 via a nonlinear transformation $\mathcal{F}(\tilde{h}_{x^2}^{*,1}, \tilde{h}_{x^2}^{*,2})$.

4 MULTI-SCALE HETEROGENEOUS GRAPH-BASED IMPUTATION NETWORK

Based on the four-phase imputation process (see Definition 8), MGIN adopts an architecture comprising four core components, as illustrated in Figure 4. Initially, the framework extracts multi-scale

temporal representations using the Heterogeneous Attribute Encoder and the Hierarchical Temporal Feature Extractor. These two components collaboratively capture hierarchical temporal features from each attribute while preserving their intrinsic attribute type characteristics. Next, the Multi-scale Dependency Mining Module constructs a multi-scale heterogeneous graph and performs missing feature imputation via a two-stage graph propagation process. Lastly, the Heterogeneous Attribute Imputation Module reconstructs missing attributes using specialized decoders tailored to each attribute type. For clarity and architectural consistency, we denote the dimensionality of hidden layers by d , noting that the actual dimensionality may vary based on the requirements of individual components.

4.1 Heterogeneous Attribute Encoder

As illustrated to the left in Figure 4, this module aims to provide representations for different attribute types (as defined in Definition 3) while preserving their intrinsic characteristics.

Spatio-temporal Attribute Encoder. Spherical coordinates (longitude λ , latitude ϕ , as defined in Definition 1 and Table 2) exhibit non-Euclidean geometry, as standard numeric encodings fail to capture spherical continuity and latitude-dependent distortion [9]. We address this by extending 3D Cartesian projections [47] with additional harmonic terms to enhance directional sensitivity. We then apply learnable affine transformations with hyperbolic tangent activations to produce embeddings \mathbf{e}_λ and \mathbf{e}_ϕ .

Timestamp τ has nested periodicities (daily, weekly, monthly, yearly cycles). To preserve these temporal patterns for neural processing, we employ multi-frequency sinusoidal encoding that captures key temporal cycles [12]. The sinusoidal outputs then pass through a linear layer with the tanh activation function, producing a d -dimensional representation $\mathbf{e}_\tau \in \mathbb{R}^d$.

Cyclical Attribute Encoder. Cyclical attributes $x_c \in \mathcal{X}_c = \{\psi, \theta\}$ exhibit boundary continuity requirements. To preserve rotational equivalence, we implement trigonometric encoding using sine and cosine transformations, followed by learnable linear layers with tanh activation. This approach maintains angular continuity in

the embedding space while ensuring smooth transitions between neighboring angular values.

Continuous Attribute Encoder. Continuous attributes $x_n \in \mathbf{X}_n = \{s, d, \ell, \beta\}$ range widely in scale, which can skew training dynamics. We address this via adaptive normalization using running statistics from the training set, with learnable scaling parameters for each attribute. A linear layer with ReLU activation then produces the final embeddings \mathbf{e}_n , helping preserve original distributions while stabilizing gradients.

Discrete Attribute Encoder. Discrete attributes $x_d \in \mathbf{X}_d = \{\eta, \chi, \kappa\}$ represent categorical semantics without intrinsic ordering. We convert them to one-hot vectors and project through learnable transformations with tanh activation to obtain embeddings \mathbf{e}_d .

Finally, all type-specific representations are concatenated into a unified embedding matrix $\mathbf{E}_t \in \mathbb{R}^{N \times d}$, preserving each attribute's intrinsic properties within a shared embedding space.

4.2 Hierarchical Temporal Feature Extractor

As illustrated in the middle-left part of Figure 4, this module captures multiple time-scale features via three complementary mechanisms: 1) Progressive temporal abstraction through increasing layer depth for finer time scales; 2) High-frequency smoothing via leaky integration in upper layers; 3) Fixed reservoir weights for computational efficiency. Concretely, we implement Deep Echo State Networks (DeepESN) [13] with leaky integrator neurons [16] as our core architecture. Compared with other sequence models (such as LSTM [14], Transformer [44], etc.), their hierarchical reservoir structure naturally produces multi-scale representations through layer-wise abstraction, while fixed random weights enable efficient processing of AIS data. Moreover, the leaky integration mechanism provides explicit control over temporal memory retention, aligning with our time-scale hierarchy in Definition 4.

For an attribute x^k at time scale $k \in \{1, \dots, 5\}$, we use $5 + 1 - k$ recurrent layers, so that higher time-scale attributes require fewer layers. The state update is governed by:

$$\begin{aligned} \mathbf{h}_t^{k-1} &= \mathbf{e}_t^k, \\ \bar{\mathbf{h}}_t^l &= \tanh\left(\mathbf{W}_h^l \cdot \mathbf{h}_t^{l-1} + \bar{\mathbf{W}}_h^{(l)} \cdot \mathbf{h}_{t-1}^l + \mathbf{b}_h^{(l)}\right), \\ \mathbf{h}_t^l &= (1 - \gamma_l) \mathbf{h}_{t-1}^l + \gamma_l \bar{\mathbf{h}}_t^l, \end{aligned} \quad (2)$$

where $\gamma_l \in (0, 1]$ is the leak rate controlling temporal retention at layer l . Random weight matrices $\mathbf{W}_h^l, \bar{\mathbf{W}}_h^l \in \mathbb{R}^{d \times d}$ and bias $\mathbf{b}_h^l \in \mathbb{R}^d$ preserve the echo state property while inducing layer-specific dynamics. Attributes at time scale k are processed exclusively by layers $l \in \{k, \dots, 5\}$, since lower-scale features ($l < k$) do not originate from higher-scale attributes (see Example 2). Equation 2 represents the forward-only version of DeepESN that enables sequential processing of streaming AIS data for real-time imputation. This formulation allows the model to impute missing values based solely on past and current information, making it suitable for safety-critical maritime applications requiring immediate response. In offline scenarios where complete trajectory data is available, we can replace the forward-only DeepESN with a **bidirectional version** to achieve enhanced performance. The default setting of MH-GIN is the bidirectional version.

Finally, we denote the consolidated multi-scale temporal representation at time step t by:

$$\mathbf{H}_t = [\mathbf{H}_t^1, \dots, \mathbf{H}_t^5], \quad (3)$$

where $\mathbf{H}_t^l = [\mathbf{h}_{t,1}^l, \dots, \mathbf{h}_{t,N_l}^l]$ aggregates temporal features at scale l , with N_l denoting both the feature count at scale l and the number of attributes from higher or equal scales ($k \geq l$).

4.3 Multi-scale Dependency Mining Module

As illustrated in the middle-right part of Figure 4, the Multi-scale Dependency Mining Module performs three main functions: 1) aligning features within same time scale for temporal consistency, 2) learning dependencies between features in partially observed data, and 3) recovering missing features from incomplete attribute sets. Conventional GNNs (e.g. GAT [45]) struggle under these conditions, as missing attributes contaminate adjacency estimation (through corrupted node similarity measures) and cause cascading errors in feature updates. Specifically, we use a two-stage propagation mechanism to capture time-varying dependencies between features at different time scales.

Stage 1: Intra-scale Propagation. This phase aligns features that occur at the same time scale to ensure their temporal patterns are synchronized.

To capture time-varying dependencies between different features at the same time scale, we model the edges within each time-scale subgraph $\mathcal{G}^k = (\mathcal{V}^k, \mathcal{E}^k)$ as dynamic rather than static. The weight of an edge is computed dynamically at each time step, explicitly modulated by corresponding features from higher time scales. Specifically, for each time scale subgraph with representations \mathbf{H}_t^k , we construct a dynamic adjacency matrix $\hat{\mathbf{A}}_t^k$ that captures time-varying dependencies between features at time scale k :

$$\hat{\mathbf{A}}^k = f_{\text{edge}}^k([\mathbf{H}^{k+1}; \dots; \mathbf{H}^5]) + \mathbf{A}^k, \quad (4)$$

where f_{edge}^k is a learnable function that computes edge weights based on both current scale features \mathbf{H}^k and higher-scale contextual features \mathbf{H}^{k+1} and \mathbf{A}^k is a bias matrix. This allows edge weights at lower scales to be adjusted based on contextual features from higher scales, enabling the model to capture time-varying dependencies. The intra-scale propagation is then performed as:

$$\tilde{\mathbf{H}}^k = (\mathbf{D}^k)^{-1/2} \cdot \hat{\mathbf{A}}^k \cdot (\mathbf{D}^k)^{-1/2} \cdot \mathbf{H}^k, \quad (5)$$

where $\mathbf{D}^k = \text{diag}(\sum_j \hat{\mathbf{A}}_{ij}^k)$ is the diagonal degree matrix that normalizes node degrees to prevent over-smoothing.

Stage 2: Cross-scale Propagation. Building on temporally aligned features, this phase integrates multi-scale features through attribute subgraphs. We keep the edges within attribute subgraphs (\mathcal{G}_x) static since higher-scale features already provide hierarchical context. Dynamic cross-scale pathways would create recursive dependencies that introduce information redundancy and hinder convergence [41]. Specifically, for each attribute x with multi-scale representations $\tilde{\mathbf{H}}_x$, we learn cross-scale interactions via:

$$\hat{\mathbf{H}}_x = (\mathbf{D}_x)^{-1/2} \cdot \hat{\mathbf{A}}_x \cdot (\mathbf{D}_x)^{-1/2} \cdot \tilde{\mathbf{H}}_x, \quad (6)$$

where the learnable matrix $\hat{\mathbf{A}}_x$ models hierarchical temporal dependencies. The final representation combines both stages through

residual connection:

$$\mathbf{h}_x^{*k} = [\mathbf{h}_x^k; \tilde{\mathbf{h}}_x^k; \hat{\mathbf{h}}_x^k]. \quad (7)$$

Theoretical Analysis. The multi-scale propagation mechanism ensures both numerical stability and robustness to input perturbations through its symmetric normalization design. We establish these properties through two key theoretical results.

Lemma 1 (Stability of Multi-scale Propagation). *The multi-scale propagation mechanism is numerically stable if the spectral radius of each propagation matrix satisfies $\rho(\mathbf{P}) = \mathbf{D}^{-1/2} \hat{\mathbf{A}} \mathbf{D}^{-1/2} \leq 1$.*

This lemma guarantees that feature representations remain bounded throughout the propagation process, preventing numerical instability such as gradient explosion or vanishing. The symmetric normalization in both intra-scale and cross-scale propagation ensures that each stage acts as a non-expansive mapping, making the complete two-stage mechanism numerically stable.

Lemma 2 (Robustness of Multi-scale Propagation). *The multi-scale propagation operator \mathcal{G} is Lipschitz continuous with bounded Lipschitz constant.*

This property ensures that small perturbations to the input features (e.g., from missing data or noise) result in predictably bounded changes in the final output. The Lipschitz continuity arises from the composition of linear operators with bounded spectral norms, making the model robust to input uncertainties commonly encountered in real-world AIS data. The proofs of these two lemmas are provided in Appendix B [23].

4.4 Heterogeneous Attribute Imputation

As illustrated in the right part of Figure 4, this module completes the imputation cycle by reconstructing missing attributes through specialized decoders for each attribute type.

Gated Fusion. The imputation process begins with a gated fusion mechanism that adaptively integrates information across multiple time scales, enabling the model to prioritize the most relevant temporal features for each missing attribute before applying type-specific reconstruction. Given the cross-scale representations $\mathbf{h}_{x^k}^*$, we apply gated fusion across time scales:

$$\mathbf{g} = \sigma(\mathbf{W}_g \cdot [\mathbf{h}_{x^k}^{*,1}; \dots; \mathbf{h}_{x^k}^{*,5}] + \mathbf{b}_g), \quad (8)$$

$$\tilde{\mathbf{e}}_{x^k} = \sum_{l=1}^5 \mathbf{g}^l \odot (\mathbf{W}_e^l \cdot \mathbf{h}_{x^k}^{*,l}), \quad (9)$$

where $\mathbf{g} \in [0, 1]^5$ are learnable gating weights, \odot denotes element-wise multiplication, $\mathbf{W}_g \in \mathbb{R}^{5 \times 5d}$ aligns dimensions for scale integration, and $\mathbf{h}_{x^k}^{*,l}$ is set to zero vector when l is lower than k .

Spatio-temporal Decoder. For coordinate attributes, we leverage the inherent spatial continuity of vessel trajectories by using local averages as base estimates, enabling the network to focus on learning incremental adjustments rather than absolute positions. Given missing coordinate attributes λ_i and ϕ_i at time step i , we first compute the mean coordinates from nearby valid observations within a window \mathcal{N}_i . The model then computes spatial adjustments δ_λ and δ_ϕ using the gated fusion embeddings $\tilde{\mathbf{e}}_{\lambda_i}, \tilde{\mathbf{e}}_{\phi_i}$ from Eq. 9. The final imputed coordinates are obtained as $\hat{\lambda} = \lambda_{\text{base}} + \delta_\lambda$ and $\hat{\phi} = \phi_{\text{base}} + \delta_\phi$.

For timestamp attribute τ , we predict time intervals between consecutive records rather than absolute timestamps, modeling temporal occurrences as a temporal point process [21, 32]. We implement a conditional intensity function $\eta(\tau)$ using neural networks with Softplus activation to ensure non-negativity. The predicted timestamp is derived as $\hat{\tau}_i = \tau_{i-1} - \frac{\log(u)}{\eta(\tau)}$, where $u \sim \text{Uniform}(0, 1)$.

Cyclical Decoder. To preserve cyclical continuity, the cyclical decoder reconstructs angular values using trigonometric projection. The decoder first transforms the gated fusion embedding $\tilde{\mathbf{e}}_{x_c}$ through a two-layer neural network to generate intermediate representation $\mathbf{h}_c \in \mathbb{R}^2$. This intermediate representation is then normalized to unit length $\hat{\mathbf{e}}_c = \frac{\mathbf{h}_c}{\|\mathbf{h}_c\|_2}$ to ensure trigonometric identity $\hat{e}_{c,1}^2 + \hat{e}_{c,2}^2 = 1$. The final angular value is decoded as $\hat{x}_c = \frac{180}{\pi} \arctan_2(\hat{e}_{c,1}, \hat{e}_{c,2})$.

Continuous Decoder. For continuous attributes, we transform the gated fusion embedding $\tilde{\mathbf{e}}_{x_n}$ through a ReLU-activated linear layer to generate normalized representation \mathbf{h}_n , then apply inverse normalization using the encoder’s parameters α, β (see Section 4.1) and training statistics μ, σ to restore the original scale as \hat{x}_n .

Discrete Decoder. For discrete attributes, we apply a linear layer followed by softmax to predict class probabilities \hat{y} from the gated fusion embedding $\tilde{\mathbf{e}}_{x_d}$, where the output dimension corresponds to the number of discrete classes $|C|$.

4.5 Training

To optimize MH-GIN, we design a comprehensive loss function that combines specialized reconstruction losses for each heterogeneous attribute type:

$$\mathcal{L}_{\text{total}} = \lambda_{\text{coo}} \mathcal{L}_{\text{coo}} + \lambda_{\tau} \mathcal{L}_{\tau} + \lambda_{\text{period}} \mathcal{L}_{\text{period}} + \lambda_{\text{cont}} \mathcal{L}_{\text{cont}} + \lambda_{\text{disc}} \mathcal{L}_{\text{disc}}, \quad (10)$$

where $\lambda_{\text{coo}}, \lambda_{\tau}, \lambda_{\text{period}}, \lambda_{\text{cont}}$, and λ_{disc} are balancing hyperparameters. The training employs specialized loss functions tailored to each attribute type: Haversine distance \mathcal{L}_{coo} [48] for spatio-temporal coordinates to capture spherical distance, mean squared error \mathcal{L}_{τ} for timestamp intervals, trigonometric encoding loss $\mathcal{L}_{\text{period}}$ for cyclical attributes to preserve periodicity, mean squared error $\mathcal{L}_{\text{cont}}$ for continuous attributes, and cross-entropy loss $\mathcal{L}_{\text{disc}}$ for discrete classification. This multi-objective design ensures that each attribute type is reconstructed according to its inherent characteristics and constraints. The more details about MH-GIN are provided in Appendix A [23].

5 EXPERIMENTS

5.1 Overall Settings

5.1.1 Datasets. We use two AIS datasets: AIS-DK from the Danish Maritime Authority [10] and AIS-US from NOAA [29]. **AIS-DK** covers Danish waters from March to December 2024, including major shipping routes in the Baltic and North Seas. It contains 18,481 vessel sequences and 10,518,249 AIS records, with an average sequence duration of 9 days. **AIS-US** spans March to May 2024, focusing on US coastal waters with dense traffic near major ports. It includes 19,091 vessel sequences and 10,546,297 AIS records, with an average sequence duration of 4 days.

Both datasets feature diverse vessel types and often capture multiple voyages per sequence, especially for short-distance operations.

Table 3: Overall effectiveness comparison on AIS-DK.

Method	Coord. (λ, ϕ)		Time τ		Head. ψ		Course θ		Speed s		Draft d		Length ℓ		Width β		Type κ
	Dist.		MAE	SMAPE	MAE	SMAPE	MAE	SMAPE	MAE	SMAPE	MAE	SMAPE	MAE	SMAPE	MAE	SMAPE	ACC
MEAN	1.313e-2	371.970	0.934	0.861	0.375	0.916	0.392	2.318	0.400	1.883	0.812	-	-	-	-	-	-
KNN	8.432e-4	34.880	0.404	0.953	0.392	0.509	0.465	0.875	0.453	1.310	0.722	-	-	-	-	-	-
Lin-ITP	5.294e-4	33.243	0.396	2.895	0.429	8.950	1.478	0.792	0.486	1.460	0.821	-	-	-	-	-	-
MF	3.314e-1	57.133	0.728	147.172	1.420	148.621	1.421	6.955	0.963	4.399	0.957	189.646	1.628	29.426	1.734	-	-
TRMF	1.341e-1	37.604	0.479	51.378	0.869	55.934	0.3972	1.580	0.381	1.968	0.839	62.843	0.333	9.308	0.325	-	-
CSDI	1.583e+0	100.375	1.588	0.388	0.390	0.432	0.246	0.745	0.397	7.786	0.625	141.077	1.514	18.240	1.120	-	-
PriSTI	1.310e+0	71.486	0.857	0.289	0.364	0.286	0.172	0.716	0.373	6.186	0.399	145.274	1.273	8.456	0.273	-	-
ImputeFormer	1.839e-1	48.375	0.689	0.486	0.434	0.776	0.386	1.344	0.492	8.713	0.782	46.055	0.194	7.687	0.259	-	-
Multi-task AIS	3.863e-4	-	-	-	-	-	-	-	-	-	-	-	-	-	-	0.322	-
PG-DPM	1.777e-4	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
MH-GIN	1.112e-4	21.413	0.127	0.049	0.036	0.141	0.071	0.198	0.252	1.017	0.126	34.192	0.178	4.896	0.175	0.669	-

Their high temporal resolution and broad coverage provide a robust basis for evaluating our methods.

5.1.2 Evaluation Metrics. Our evaluation metrics are tailored to the distinct characteristics of AIS data attributes. For continuous numerical values, we employ Mean Absolute Error (MAE) and Symmetric Mean Absolute Percentage Error (SMAPE). For categorical attributes, we utilize Accuracy (ACC).

5.1.3 Baseline Methods. To evaluate the performance of our proposed method, we compare with classic models and state-of-the-art methods for multi-variable time-series imputation, spatio-temporal imputation, and trajectory imputation. The baselines include: statistical methods (MEAN, KNN, Lin-ITP), classic models (MF, TRMF [54]), multi-variable time-series imputation (CSDI [40]), spatiotemporal imputation (ImputeFormer, PriSTI [24]), and trajectory imputation (Multi-task AIS [30], PG-DPM [55]).

5.1.4 Masking Strategies. AIS data presents distinctive missing patterns across various time scales (as discussed in Section 3.3). To accurately simulate these real-world scenarios, we implement three targeted masking strategies with mask ratio r :

- **Point Masking** (Scales 1 and 2): Randomly masks individual values throughout the sequence with probability r , simulating sporadic missingness characteristic of high-frequency attributes.
- **Block Masking** (Scales 3 and 4): Masks continuous segments within individual voyages throughout the sequence with probability r , capturing voyage-phase-related missing patterns common to these attributes.
- **Entire Masking** (Scale 5): Masks entire attribute sequences for a proportion r of vessels, replicating the systematic absence of vessel-specific data frequently observed in AIS records.

5.1.5 Noise Injection Strategies. To evaluate the robustness of our proposed method, we inject noise into the AIS data to simulate real-world data corruption scenarios. We design a unified noise injection strategy controlled by parameter γ that provides comparable corruption levels across all heterogeneous attribute types. For continuous attributes, we add Gaussian noise scaled by the attribute's own value $x'_n = x_n + \mathcal{N}(0, (\gamma x_n)^2)$. For spatio-temporal coordinates, noise is scaled by movement dynamics to simulate GPS drift $\lambda' = \lambda + \mathcal{N}(0, (\gamma \sigma_{\lambda})^2)$ and $\phi' = \phi + \mathcal{N}(0, (\gamma \sigma_{\Delta\phi})^2)$. For

timestamps, we corrupt time intervals rather than absolute values $\Delta\tau'_i = \Delta\tau_i + \max(0, \mathcal{N}(0, (\gamma \cdot \sigma_{\Delta\tau})^2))$. For cyclical attributes like headings, we add angular noise with proper wrap-around handling $x'_c = (x_c + \mathcal{N}(0, (\gamma \cdot \sigma_{\Delta x_c})^2)) \pmod{360}$. For discrete categories, we implement label flipping with probability γ . This unified approach ensures that γ serves as an interpretable control parameter for systematic noise injection across the heterogeneous dataset.

5.1.6 Experimental Settings. All experiments are conducted on a server with Intel Xeon Processor (Icelake) CPUs, 100GB RAM, and two NVIDIA A10 GPUs (each with 23GB memory). For training our model, we use the Adam optimizer with a learning rate of $1e^{-3}$ and weight decay of $1e^{-4}$. The batch size is set to 64, and we train for 100 epochs. Early stopping with a patience of 10 epochs was applied to prevent overfitting.

For data processing, we use an 80%–10%–10% split for training, validation, and testing, respectively. For the masked value generation, we create synthetic missing data according to the three strategies described earlier, with mask ratios of 10%, 20%, 30%, 40% and 50% to evaluate the robustness of different imputation methods. For the noise injection, we conduct experiments by varying the noise intensity γ progressively from 0.0 to 0.5. The default mask ratio is set to 30% for all experiments unless otherwise specified.

For methods, MEAN, Lin-ITP, and MF, no hyperparameters need to be set. We set k of KNN to 20. For TRMF, CSDI, PriSTI, ImputeFormer, and Multi-task AIS, we use the default settings provided by the authors. The more details about the experiments are provided in Appendix C [23].

5.2 Performance Evaluation

5.2.1 Effectiveness Analysis. Tables 3 and 4 present the performance comparison between MH-GIN and baseline methods. Navigation status η and hazardous cargo type χ are excluded from Tables 3 and 4 due to the lack of suitable baseline methods capable of imputing these specific categorical attributes. However, we emphasize that both categorical attributes are fully included in all subsequent experiments, including ablation studies, missing ratio analysis, and robustness evaluation.

MH-GIN consistently outperforms all baseline methods across both datasets and all attribute types, demonstrating its strong capability to leverage multi-scale dependencies between attributes for

Table 4: Overall effectiveness comparison on AIS-US.

Method	Coord. (λ, ϕ)		Time τ		Head. ψ		Course θ		Speed s		Draft d		Length ℓ		Width β		Type κ
	Dist.		MAE	SMAPE	MAE	SMAPE	MAE	SMAPE	MAE	SMAPE	MAE	SMAPE	MAE	SMAPE	MAE	SMAPE	ACC
MEAN	1.604e-2	260.317	0.491		0.908	0.370	1.203	0.449	2.433	1.344	3.164	0.476	-	-	-	-	-
KNN	7.818e-4	44.358	0.516		0.452	0.391	0.846	0.640	0.702	1.020	2.134	0.322	-	-	-	-	-
Lin-ITP	4.896e-4	31.013	0.486		7.617	0.365	28.444	0.487	0.687	0.968	2.429	0.394	-	-	-	-	-
MF	6.830e-1	160.671	0.702		128.235	1.115	131.286	1.008	3.701	1.303	5.499	1.123	181.249	1.763	29.730	1.842	-
TRMF	3.536e-1	232.299	0.945		68.246	0.491	82.993	0.580	1.489	1.279	2.617	0.462	58.320	0.366	8.326	0.307	-
CSDI	1.552e+0	373.14	1.993		0.437	0.245	0.689	0.364	0.663	0.857	8.096	1.745	154.77	1.887	18.047	1.025	-
PriSTI	1.663e+0	329.957	1.608		0.348	0.233	0.595	0.329	0.591	0.787	4.997	0.510	134.655	1.202	10.038	0.397	-
ImputeFormer	2.104e-1	110.320	0.536		1.138	0.570	1.216	0.505	1.865	0.976	4.360	0.943	45.337	0.318	6.820	0.268	-
Multi-task AIS	3.627e-4	-	-		-	-	-	-	-	-	-	-	-	-	-	-	0.355
PG-DPM	1.512e-4	-	-		-	-	-	-	-	-	-	-	-	-	-	-	-
MH-GIN	8.281e-5	15.146	0.016		0.104	0.072	0.461	0.184	0.368	1.182	1.899	0.314	38.263	0.273	5.382	0.213	0.446

accurate AIS data imputation. Specifically, we observe the following: 1) Spatio-temporal attributes – Coordinate (λ, ϕ) and Timestamp τ : Lin-ITP achieves the best timestamp imputation performance by ensuring interpolated values fall within the temporal range of consecutive records. Similarly, PG-DPM performs well in coordinate imputation by incorporating physics-guided constraints on vessel movement. Although these methods are specifically designed for spatio-temporal attributes, MH-GIN still surpasses both with 35.59%–96.71% performance gains, which strongly demonstrates the effectiveness of modeling multi-scale dependencies. It also confirms that MH-GIN’s spatio-temporal encoder-decoder is not only effective in its own right, but also sufficiently expressive to support subsequent multi-scale dependency mining. 2) Cyclical attributes – Heading ψ and Course θ : Among baselines, PriSTI achieves the best results for cyclical attribute imputation. However, due to the dynamic nature of spatial information in AIS data, PriSTI and ImputeFormer fail to fully utilize spatial context, resulting in comparable performance to CSDI. While these methods outperform classic baselines, they struggle with boundary continuity in cyclical attributes. In contrast, MH-GIN achieves 22.52%–98.93% improvements over PriSTI, demonstrating the strength of its cyclical encoder-decoder in modeling cyclical patterns while preserving boundary continuity. 3) Continuous attributes – Speed s , Draft d , Length ℓ , and Width β : MH-GIN yields 2.48%–72.35% improvements over PriSTI and ImputeFormer, primarily due to its ability to handle scale variation through a dedicated continuous encoder-decoder architecture. Notably, CSDI, PriSTI and ImputeFormer generally assume that all variables are continuous attributes, which gives them a natural advantage when dealing with this attribute type. Despite this, MH-GIN still surpasses their performance, highlighting the importance and effectiveness of mining multi-scale dependencies across heterogeneous attributes for more accurate imputation. 4) For discrete attribute imputation (Type κ), MH-GIN reduces error rates by 14.11%–51.18% compared to Multi-task AIS. This superior performance stems from MH-GIN’s comprehensive modeling approach that incorporates all 12 attributes across different time scales, whereas Multi-task AIS only utilizes limited information (spatio-temporal attributes and vessel type). By effectively capturing cross-scale dependencies among the heterogeneous attributes,

Table 5: Overall time cost and memory cost comparison.

Method	AIS-DK		AIS-US	
	Time (s)	Memory (MB)	Time (s)	Memory (MB)
MEAN	1.24	-	1.28	-
KNN	5.39	-	7.50	-
Lin-ITP	1.26	-	1.29	-
MF	27.24	-	28.29	-
TRMF	34.02	-	35.86	-
CSDI	60.48	1.72	60.74	1.92
PriSTI	94.92	2.89	100.80	3.19
ImputeFormer	14.10	5.13	15.42	5.52
Multi-task AIS	38.30	1.66	40.59	1.84
PG-DPM	120.46	1.79	128.32	1.99
MH-GIN	10.9	1.43	11.2	1.83

MH-GIN establishes more robust correlations that significantly enhance vessel type imputation accuracy.

5.2.2 Efficiency Analysis. Table 5 presents a comparison of computational efficiency in terms of inference time on the test dataset and memory consumption (model size) for each method. Key observations include: 1) Statistical methods (MEAN, KNN, Lin-ITP) achieve significantly lower computational times than parametric approaches. 2) MH-GIN demonstrates superior efficiency among neural methods, with inference times of 10.9s and 11.2s (22.7% and 27.4% faster than ImputeFormer) on AIS-DK and AIS-US datasets. 3) MH-GIN maintains minimal memory costs at 1.47MB and 1.92 MB. These efficiency gains come from the Multi-scale Dependency Mining Module’s efficient and simple architecture that eliminates complex operations required by complex neural networks.

5.3 Ablation Study

To evaluate the contribution of each component of MH-GIN, we conducted a comprehensive ablation study by removing or replacing key components and analyzing the impact of different attribute types. Table 6 presents the results on model components.

Table 6: Effectiveness analysis of model components.

Model Variant	Coord. (λ, ϕ)		Time τ		Head. ψ		Course θ		Speed s		Nav. η	Cargo χ	Draft d		Length ℓ		Width β	Type κ	
	Dist.		MAE	SMAPE	MAE	SMAPE	MAE	SMAPE	MAE	SMAPE	ACC	ACC	MAE	SMAPE	MAE	SMAPE	ACC		
AIS-DK	w/o Spa	1.364e-3	45.326	0.465	0.072	0.046	0.171	0.085	0.242	0.283	0.802	0.718	1.323	0.174	42.329	0.223	6.258	0.216	0.584
	w/o Cyc	1.316e-4	25.769	0.158	0.062	0.041	0.176	0.269	0.216	0.261	0.881	0.769	1.134	0.143	35.106	0.188	4.994	0.184	0.651
	w/o Con	1.151e-4	22.312	0.138	0.059	0.041	0.147	0.079	0.211	0.261	0.882	0.771	1.126	0.139	35.364	0.216	5.124	0.196	0.659
	w/o Dis	1.249e-4	22.631	0.141	0.059	0.042	0.156	0.088	0.238	0.269	-	-	1.159	0.152	35.399	0.189	5.043	0.199	-
	w/o M2-B	1.446e-4	28.643	0.179	0.069	0.053	0.192	0.106	0.246	0.286	0.813	0.625	1.249	0.171	38.352	0.204	5.984	0.215	0.532
	w/o M3-S1	1.530e-4	31.825	0.372	0.872	0.456	0.185	0.102	0.758	0.410	0.322	0.380	1.385	0.765	-	-	-	-	-
	w/o M3-S1-D	1.117e-4	21.652	0.129	0.050	0.037	0.143	0.074	0.203	0.256	0.898	0.792	1.019	0.127	34.388	0.180	4.901	0.178	0.665
	w/o M3-S2	1.256e-4	23.125	0.142	0.051	0.038	0.145	0.076	0.206	0.257	0.890	0.780	1.095	0.135	34.825	0.183	4.950	0.180	0.658
	MH-GIN	1.112e-4	21.413	0.127	0.049	0.036	0.141	0.071	0.198	0.252	0.904	0.798	1.017	0.126	34.192	0.178	4.896	0.175	0.669
AIS-US	w/o Spa	1.465e-4	28.59	0.082	0.132	0.125	0.493	0.239	0.469	1.326	0.592	0.524	2.258	0.349	41.536	0.329	5.869	0.373	0.394
	w/o Cyc	9.242e-5	17.511	0.038	0.127	0.093	0.513	0.243	0.379	1.196	0.631	0.542	1.931	0.324	38.996	0.281	5.412	0.219	0.426
	w/o Con	8.624e-5	15.866	0.018	0.119	0.085	0.479	0.189	0.382	1.191	0.633	0.545	1.965	0.347	39.26	0.284	5.762	0.373	0.419
	w/o Dis	8.963e-5	15.976	0.021	0.127	0.082	0.483	0.189	0.386	1.196	-	-	1.953	0.367	39.104	0.292	5.536	0.227	-
	w/o M2-B	1.265e-4	17.937	0.025	0.134	0.092	0.579	0.231	0.439	1.368	0.582	0.446	2.218	0.374	41.726	0.314	5.936	0.252	0.364
	w/o M3-S1	1.372e-4	26.325	0.061	1.142	0.412	1.590	0.538	1.252	1.178	0.175	0.128	2.240	0.375	-	-	-	-	-
	w/o M3-S1-D	8.364e-5	15.259	0.017	0.108	0.075	0.467	0.186	0.370	1.184	0.644	0.555	1.910	0.319	38.681	0.279	5.388	0.215	0.433
	w/o M3-S2	8.731e-5	15.553	0.018	0.109	0.076	0.468	0.187	0.372	1.186	0.641	0.552	1.918	0.320	38.811	0.231	5.391	0.224	0.431
	MH-GIN	8.281e-5	15.146	0.016	0.104	0.072	0.461	0.184	0.368	1.182	0.651	0.559	1.899	0.314	38.263	0.273	5.382	0.213	0.446

Table 7: Robustness evaluation under different missing ratios.

Missing Rate	Coord. (λ, ϕ)		Time τ		Head. ψ		Course θ		Speed s		Nav. η	Cargo χ	Draft d		Length ℓ		Width β	Type κ	
	Dist.		MAE	SMAPE	MAE	SMAPE	MAE	SMAPE	MAE	SMAPE	ACC	ACC	MAE	SMAPE	MAE	SMAPE	ACC		
AIS-DK	10%	0.979e-5	19.319	0.114	0.045	0.033	0.136	0.070	0.197	0.247	0.917	0.814	0.883	0.121	32.951	0.173	4.536	0.165	0.689
	20%	1.043e-4	20.475	0.121	0.046	0.034	0.139	0.071	0.198	0.249	0.908	0.805	0.901	0.124	33.674	0.176	4.717	0.170	0.682
	30%	1.112e-4	21.413	0.127	0.049	0.036	0.141	0.071	0.198	0.252	0.904	0.798	1.017	0.126	34.192	0.178	4.896	0.175	0.669
	40%	3.229e-4	23.121	0.133	0.051	0.037	0.146	0.075	0.206	0.261	0.887	0.717	1.196	0.141	35.462	0.184	5.096	0.183	0.652
	50%	1.383e-3	24.868	0.147	0.055	0.041	0.151	0.078	0.213	0.267	0.868	0.639	1.453	0.152	36.842	0.191	5.319	0.191	0.628
AIS-US	10%	7.119e-5	14.313	0.014	0.101	0.069	0.451	0.177	0.356	1.164	0.664	0.573	1.838	0.306	37.536	0.271	5.239	0.206	0.456
	20%	7.776e-5	14.822	0.015	0.103	0.072	0.456	0.181	0.362	1.174	0.659	0.563	1.875	0.313	38.109	0.273	5.304	0.210	0.449
	30%	8.281e-5	15.146	0.016	0.104	0.072	0.461	0.184	0.368	1.182	0.651	0.559	1.899	0.314	38.263	0.273	5.382	0.213	0.446
	40%	4.113e-4	15.861	0.017	0.110	0.077	0.473	0.189	0.372	1.192	0.634	0.546	1.962	0.324	39.246	0.283	5.469	0.219	0.424
	50%	9.019e-4	16.533	0.018	0.115	0.081	0.484	0.194	0.385	1.206	0.621	0.529	2.023	0.333	40.115	0.291	5.579	0.225	0.411

We have following observations: 1) w/o Spa, w/o Cyc, w/o Con, w/o Dis: Removing type-specific encoders and decoders for spatio-temporal, cyclical, continuous, and discrete attributes leads to significant performance degradation in their corresponding attribute types. Performance also drops for other attribute types due to missing well-encoded representations, demonstrating the interdependencies between heterogeneous attributes. 2) w/o M2-B: Replacing bidirectional DeepESN with forward-only variant reduces performance across all attributes. Nevertheless, M2-B still outperforms the baseline methods (as shown in Table 3 and Table 4), demonstrating the effectiveness of MH-GIN in real-time scenarios. 3) w/o M3-S1: Removing intra-scale propagation severely impacts performance across all attributes, demonstrating that MH-GIN cannot effectively capture inter-attribute dependencies without this component. This highlights the critical importance of modeling complementary relationships between heterogeneous attributes within each time scale. 4) w/o M3-S1-D: Removing dynamic edges within time-scale subgraphs while keeping static graph structure shows performance

degradation compared to the full model, demonstrating that dynamic relationship modeling is essential for capturing the evolving dependencies between attributes at each time scale. 5) w/o M3-S2: Eliminating cross-scale propagation results in performance degradation, particularly for attributes with dependencies across scales. This demonstrates that cross-scale relationships are essential for effectively leveraging complementary information across different temporal dynamics.

5.4 Robustness of MH-GIN

Performance under Different Missing Ratios. We analyze MH-GIN’s robustness under varying missing ratios (10%-50%) in Table 7. The results yield the following observations: 1) Performance degrades gradually with increasing missing ratios across all attributes, indicating model stability under data sparsity. 2) Non-coordinate attributes exhibit near-linear degradation without catastrophic collapse. The hierarchical structure enables lower-scale attributes to

Table 8: Robustness evaluation under noisy conditions.

Noisy Intensity	Coord. (λ, ϕ)		Time τ		Head. ψ		Course θ		Speed s		Nav. η	Cargo χ	Draft d		Length ℓ		Width β		Type κ
	Dist.		MAE	SMAPE	MAE	SMAPE	MAE	SMAPE	MAE	SMAPE	ACC	ACC	MAE	SMAPE	MAE	SMAPE	MAE	SMAPE	ACC
AIS-DK	0.00	1.112e-4	21.413	0.127	0.049	0.036	0.141	0.071	0.198	0.252	0.904	0.798	1.017	0.126	34.192	0.178	4.896	0.175	0.669
	0.01	1.123e-4	22.017	0.129	0.049	0.038	0.143	0.073	0.203	0.253	0.902	0.798	1.056	1.129	34.384	0.180	4.912	0.178	0.667
	0.015	1.366e-4	22.656	0.130	0.051	0.039	0.144	0.076	0.205	0.256	0.899	0.796	1.213	1.132	34.681	0.181	4.947	0.179	0.665
	0.02	1.647e-4	22.931	0.132	0.053	0.041	0.146	0.078	0.208	0.258	0.898	0.795	1.239	1.136	35.152	0.183	4.973	0.182	0.662
	0.025	2.148e-4	23.172	0.133	0.054	0.043	0.149	0.081	0.210	0.261	0.897	0.792	1.305	1.142	35.503	0.186	5.062	0.185	0.658
AIS-US	0.00	8.281e-5	15.146	0.016	0.104	0.072	0.461	0.184	0.368	1.182	0.651	0.559	1.899	0.314	38.263	0.273	5.382	0.213	0.446
	0.01	8.316e-5	15.265	0.016	0.106	0.074	0.463	0.188	0.369	1.184	0.648	0.557	1.912	0.316	38.426	0.275	5.399	0.214	0.445
	0.015	8.999e-5	15.382	0.016	0.107	0.076	0.466	0.190	0.372	1.188	0.645	0.554	1.946	0.319	38.873	0.278	5.402	0.217	0.443
	0.02	9.642e-5	15.503	0.017	0.109	0.079	0.468	0.192	0.375	1.191	0.643	0.552	1.980	0.321	39.143	0.281	5.436	0.219	0.441
	0.025	1.017e-4	15.621	0.017	0.112	0.081	0.472	0.195	0.379	1.195	0.640	0.549	2.014	0.325	39.479	0.284	5.475	0.222	0.438

Table 9: Effectiveness of MH-GIN in scenario where multiple attributes are missing at the same time.

Missing Number	Coord. (λ, ϕ)		Time τ		Head. ψ		Course θ		Speed s		Nav. η	Cargo χ	Draft d		Length ℓ		Width β		Type κ
	Dist.		MAE	SMAPE	MAE	SMAPE	MAE	SMAPE	MAE	SMAPE	ACC	ACC	MAE	SMAPE	MAE	SMAPE	MAE	SMAPE	ACC
AIS-DK	1-2	1.038e-4	21.316	0.113	0.045	0.031	0.132	0.061	0.193	0.241	0.916	0.807	0.987	0.119	33.992	0.162	4.632	0.168	0.681
	3-4	1.108e-4	21.388	0.125	0.047	0.034	0.139	0.067	0.196	0.246	0.911	0.802	1.004	0.123	34.014	0.167	4.784	0.171	0.673
	5-6	1.726e-4	22.632	0.139	0.056	0.044	0.146	0.078	0.207	0.259	0.892	0.784	1.126	0.141	34.369	0.183	4.996	0.185	0.661
AIS-US	1-2	8.103e-5	14.936	0.013	0.091	0.065	0.448	0.178	0.356	1.169	0.661	0.569	1.821	0.298	38.184	0.253	5.324	0.199	0.454
	3-4	8.226e-5	15.012	0.015	0.096	0.070	0.452	0.181	0.359	1.173	0.657	0.562	1.832	0.302	38.226	0.261	5.371	0.208	0.449
	5-6	8.394e-5	15.233	0.019	0.116	0.078	0.477	0.189	0.371	1.189	0.648	0.554	1.904	0.323	38.284	0.276	5.394	0.226	0.441

generate higher-scale features, allowing the model to leverage abundant high-scale features for reasonable imputation performance. 3) Coordinate attributes (lowest time scale) show significant performance drops at high missing ratios due to insufficient low-scale features and limited constraints from coarse high-scale features. However, such extreme sparsity is rare in practice, where coordinates typically remain densely observed while non-spatio-temporal attributes exhibit higher missing ratios. MH-GIN maintains strong performance under high missing ratios, demonstrating robustness in practical settings.

Performance under Different Noise Intensities. To evaluate the robustness of MH-GIN under noisy conditions, we inject noise into the data with progressively varying noise intensity γ to simulate realistic corruption scenarios. As shown in Table 8, MH-GIN exhibits remarkable resilience across noise intensities ranging from 0 to 0.025. The model exhibits graceful performance degradation in a near-linear fashion without catastrophic failure, which can be attributed to the multi-scale architecture’s capability to leverage contextual information from less corrupted attributes and scales, thereby effectively mitigating the impact of local perturbations.

Performance under Multiple Attributes Missing Scenarios. To evaluate MH-GIN’s robustness under simultaneous multi-attribute missing scenarios, we assess the model’s performance as the number of missing attributes increases. As shown in Table 9, while MH-GIN’s performance degrades as the number of missing attributes increases, the degradation remains moderate and controlled. Specifically, even with up to six missing attributes, MH-GIN maintains stable accuracy across both numerical and categorical features. This

shows that MH-GIN effectively exploits remaining correlations to stay stable under severe missing conditions.

6 CONCLUSION AND FUTURE WORK

Our study highlights the importance of modeling multi-scale dependencies between heterogeneous attributes for missing attribute imputation in AIS data. We introduce MH-GIN, a Multi-scale Heterogeneous Graph-based Imputation Network for AIS data. MH-GIN is the first method that can effectively model multi-scale dependencies between heterogeneous attributes for missing attribute imputation in AIS data. MH-GIN first captures the multi-scale temporal features of each attribute while preserving their intrinsic properties. It then constructs a multi-scale heterogeneous graph to model the complex dependencies between attributes, enabling effective missing attribute imputation through two-stage graph propagation. Experiments on two real-world AIS datasets demonstrate MH-GIN’s superior performance over state-of-the-art methods across all attribute types, achieving an average 57% reduction in imputation error, while maintaining computational efficiency.

In future work, we plan to extend MH-GIN to other domains with heterogeneous attributes across different time scales.

ACKNOWLEDGMENTS

The research leading to the results presented in this paper has received funding from the European Union’s Horizon Europe programme through the MobiSpaces project (grant agreement no. 101070279). Tianyi Li is the corresponding author.

REFERENCES

[1] Edgar Acuna and Caroline Rodriguez. 2004. The treatment of missing values and its effect on classifier accuracy. In *Meeting of the International Federation of Classification Societies*. 639–647.

[2] Mehran Amiri and Richard Jensen. 2016. Missing data imputation using fuzzy-rough methods. *Neurocomput.* 205 (2016), 152–164.

[3] Asian Development Bank. 2023. Methodological Framework for Unlocking Maritime Insights Using Automatic Identification System Data: A Special Supplement of Key Indicators for Asia and the Pacific.

[4] Parikshit Bansal, Prathamesh Deshpande, and Sunita Sarawagi. 2021. Missing Value Imputation on Multidimensional Time Series. *Proc. VLDB Endow.* 14, 11 (2021), 2533–2545.

[5] Gustavo EAPA Batista, Maria Carolina Monard, et al. 2002. A study of K-nearest neighbour as an imputation method. *His* 87, 251–260 (2002), 48.

[6] Wei Cao, Dong Wang, Jian Li, Hao Zhou, Lei Li, and Yitan Li. 2018. BRITS: Bidirectional Recurrent Imputation for Time Series. In *NeurIPS*. 6776–6786.

[7] Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, David Sontag, and Yan Liu. 2018. Recurrent Neural Networks for Multivariate Time Series with Missing Values. *Scientific Reports* 8, 1 (2018), 6085.

[8] Andrea Cini, Ivan Marisca, and Cesare Alippi. 2022. Filling the G_ap_s: Multivariate Time Series Imputation by Graph Neural Networks. In *ICLR*.

[9] Taco S. Cohen, Mario Geiger, Jonas Köhler, and Max Welling. 2018. Spherical CNNs. In *ICLR*.

[10] Danish Maritime Authority. 2024. AIS-DK (Last access: 2025.10). <https://aisdata.ais.dk/?prefix=2024/>.

[11] Wenjie Du, David Côté, and Yan Liu. 2023. SAITS: Self-attention-based imputation for time series. *Expert Syst. Appl.* 219 (2023), 119619.

[12] Yutong Feng, Qiongyan Wang, Yutong Xia, Junlin Huang, Siru Zhong, and Yuxuan Liang. 2024. Spatio-temporal field neural networks for air quality inference. In *IJCAI*. 7260–7268.

[13] Claudio Gallicchio and Alessio Micheli. 2017. Deep Echo State Network (Deep-ESN): A Brief Survey. *CoRR* abs/1712.04323.

[14] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Comput.* 9, 8 (1997), 1735–1780.

[15] IMO. 2013. Long-range identification and tracking (Last access: 2025.10). <https://www.imo.org/en/ourwork/safety/pages/lrit.aspx>.

[16] Herbert Jaeger, Mantas Lukosevicius, Dan Popovici, and Udo Sievert. 2007. Optimization and applications of echo state networks with leaky-integrator neurons. *Neural Networks* 20, 3 (2007), 335–352.

[17] Mourad Khayati, Alberto Lerner, Zakhari Tymchenko, and Philippe Cudré-Mauroux. 2020. Mind the Gap: An Experimental Evaluation of Imputation of Missing Values Techniques in Time Series. *Proc. VLDB Endow.* 13, 5 (2020), 768–782.

[18] Tianyi Li, Lu Chen, Christian S Jensen, and Torben Bach Pedersen. 2021. TRACE: Real-time compression of streaming trajectories in road networks. *Proc. VLDB Endow.* 14, 7 (2021), 1175–1187.

[19] Tianyi Li, Lu Chen, Christian S Jensen, Torben Bach Pedersen, Yunjun Gao, and Jilin Hu. 2022. Evolutionary clustering of moving objects. In *ICDE*. 2399–2411.

[20] Tianyi Li, Ruikai Huang, Lu Chen, Christian S Jensen, and Torben Bach Pedersen. 2020. Compression of uncertain trajectories in road networks. *Proc. VLDB Endow.* 13, 7 (2020), 1050–1063.

[21] Haitao Lin, Lirong Wu, Guojiang Zhao, Pai Liu, and Stan Z. Li. 2022. Exploring Generative Neural Temporal Point Process. *Trans. Mach. Learn. Res.* 2022 (2022).

[22] Zachary C. Lipton, David C. Kale, and Randall C. Wetzel. 2016. Directly Modeling Missing Data in Sequences with RNNs: Improved Classification of Clinical Time Series. In *MLHC*, Vol. 56. 253–270.

[23] Hengyu Liu, Tianyi Li, Yuqiang He, Kristian Torp, Yushuai Li, and Christian S. Jensen. 2025. MH-GIN: Multi-scale Heterogeneous Graph-based Imputation Network for AIS Data (Extended Version). *CoRR* abs/2507.20362.

[24] Mingzhe Liu, Han Huang, Hao Feng, Leilei Sun, Bowen Du, and Yanjie Fu. 2023. PriSTI: A Conditional Diffusion Framework for Spatiotemporal Imputation. In *ICDE*. 1927–1939.

[25] Piotr Lubkowski and Dariusz Laskowski. 2017. Assessment of Quality of Identification of Data in Systems of Automatic Licence Plate Recognition. In *TST*, Vol. 715. 482–493.

[26] Ivan Marisca, Andrea Cini, and Cesare Alippi. 2022. Learning to Reconstruct Missing Data from Spatiotemporal Graphs with Sparse Observations. In *NeurIPS*, Vol. 35. 32069–32082.

[27] Patrick E McKnight, Katherine M McKnight, Souraya Sidani, and Aurelio Jose Figueiredo. 2007. *Missing data: A gentle introduction*. Guilford Press.

[28] Mashaal Musleh and Mohamed F Mokbel. 2023. Kamel: A Scalable BERT-based System for Trajectory Imputation. *Proc. VLDB Endow.* 17, 3 (2023), 523–538.

[29] National Oceanic and Atmospheric Administration. 2024. AIS-US (Last access: 2025.10). <https://coast.noaa.gov/htdata/CMSP/AISDataHandler/2024/index.html>.

[30] Duong Nguyen, Rodolphe Vadaine, Guillaume Hajduch, René Garello, and Ronan Fablet. 2018. A Multi-Task Deep Learning Architecture for Maritime Surveillance Using AIS Data Streams. In *DSAA*. 331–340.

[31] Tong Nie, Guoyang Qin, Wei Ma, Yuewen Mei, and Jian Sun. 2024. ImputeFormer: Low Rankness-Induced Transformers for Generalizable Spatiotemporal Imputation. In *KDD*. 2260–2271.

[32] Takahiro Omi, Naonori Ueda, and Kazuyuki Aihara. 2019. Fully Neural Network based Model for General Temporal Point Processes. In *NeurIPS*. 2120–2129.

[33] Mengshi Qi, Jie Qin, Yu Wu, and Yi Yang. 2020. Imitative Non-Autoregressive Modeling for Trajectory Forecasting and Imputation. In *CVPR*. 12733–12742.

[34] Yichen Shi, Hongye Gao, and Weixiong Rao. 2023. TIGAN: Trajectory Imputation via Generative Adversarial Network. In *ADMA* (5), Vol. 14180. 195–209.

[35] Junjun Si, Jin Yang, Yang Xiang, Hanqiu Wang, Li Li, Rongqing Zhang, Bo Tu, and Xiangqun Chen. 2024. TrajBERT: BERT-Based Trajectory Recovery With Spatial-Temporal Refinement for Implicit Sparse Trajectories. *IEEE Trans. Mob. Comput.* 23, 5 (2024), 4849–4860.

[36] Shaoxu Song, Yu Sun, Aoqian Zhang, Lei Chen, and Jianmin Wang. 2020. Enriching Data Imputation under Similarity Rule Constraints. *IEEE Trans. Knowl. Data Eng.* 32, 2 (2020), 275–287.

[37] Yumeng Song, Yu Gu, Tianyi Li, Yushuai Li, Christian S Jensen, and Ge Yu. 2024. Quantifying Point Contributions: A Lightweight Framework for Efficient and Effective Query-Driven Trajectory Simplification. *Proc. VLDB Endow.* 18, 2 (2024), 453–465.

[38] Spire. 2024. Not all AIS data is equal: How to identify and ensure AIS data quality (Last access: 2025.10). <https://spire.com/blog/maritime/not-all-ais-data-is-equal-how-to-identify-and-ensure-ais-data-quality/>.

[39] Raphael Sturgis, Valentin Emiya, Basile Courtois, and Pierre Garreau. 2022. Vessel Behaviour Classification from AIS Without Geographical Biases. In *ITSC*. 2465–2470.

[40] Yusuke Tashiro, Jiaming Song, Yang Song, and Stefano Ermon. 2021. CSDI: Conditional Score-based Diffusion Models for Probabilistic Time Series Imputation. In *NeurIPS*. 24804–24816.

[41] Naftali Tishby and Noga Zaslavsky. 2015. Deep learning and the information bottleneck principle. In *IEEE ITW*. 1–5.

[42] U.S. Coast Guard. 2017. User's Manual of Vessel Traffic Service Prince William Sound Area (Last access: 2025.10). <https://www.commerce.alaska.gov/web/Portals/5/pub/2017VTSPWSUseRManual12thEdition.pdf>.

[43] U.S. Coast Guard. 2024. User's Manual of US Coast Guard Vessel Traffic Service Puget Sound (Last access: 2025.10). [https://www.navcen.uscg.gov/sites/default/files/pdf/VT%20User%20Guides/VTS_PS_UsersManual_\(2024\).pdf#page=4.09](https://www.navcen.uscg.gov/sites/default/files/pdf/VT%20User%20Guides/VTS_PS_UsersManual_(2024).pdf#page=4.09).

[44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NeurIPS*, Vol. 30.

[45] Petar Veličković, Guillelmo Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. In *ICLR (Poster)*.

[46] Vizion. 2025. How to Assess and Ensure High-Quality AIS Data (Last access: 2025.10). <https://www.vizionapi.com/blog/assess-ensure-high-quality-ais-data>.

[47] Wikipedia contributors. 2024. Cartesian Coordinate System (Last access: 2025.10). https://en.wikipedia.org/wiki/Cartesian_coordinate_system.

[48] Wikipedia contributors. 2024. Haversine Formula (Last access: 2025.10). https://en.wikipedia.org/wiki/Haversine_formula.

[49] Werner Wothke. 2000. Longitudinal and multilevel modeling with missing data. In *Modeling Longitudinal and Multilevel Data*. Psychology Press, 205–224.

[50] Junyang Wu, Tianyi Li, Lu Chen, Yunjun Gao, and Ziheng Wei. 2023. SEA: a scalable entity alignment system. In *SIGIR*. 3175–3179.

[51] Zhiye Xie, Enmei Tu, Xianping Fu, Guoliang Yuan, and Yi Han. 2025. AIS Data-Driven Maritime Monitoring Based on Transformer: A Comprehensive Review. *CoRR* abs/2505.07374.

[52] Yi Xu, Armin Bazarjani, Hyung-Gun Chi, Chiho Choi, and Yun Fu. 2023. Uncovering the Missing Pattern: Unified Framework Towards Trajectory Imputation and Prediction. In *CVPR*. 9632–9643.

[53] Jinsung Yoon, William R. Zame, and Mihaela van der Schaar. 2019. Estimating Missing Data in Temporal Data Streams Using Multi-Directional Recurrent Neural Networks. *IEEE Trans. Biomed. Eng.* 66, 5 (2019), 1477–1490.

[54] Hsiang-Fu Yu, Nikhil Rao, and Inderjit S. Dhillon. 2016. Temporal Regularized Matrix Factorization for High-dimensional Time Series Prediction. In *NeurIPS*. 847–855.

[55] Zhiwen Zhang, Zipei Fan, Zewu Lv, Xuan Song, and Ryosuke Shibasaki. 2024. Long-Term Vessel Trajectory Imputation with Physics-Guided Diffusion Probabilistic Model. In *KDD*. 4398–4407.

[56] Cong Zhao, Andi Song, Yuchuan Du, and Biao Yang. 2022. TrajGAT: A map-embedded graph attention network for real-time vehicle trajectory imputation of roadside perception. *Transportation research part C: emerging technologies* 142 (2022), 103787.