# AISFlow: Boundary-Informed Flow Matching for Long-Term AIS Trajectory Imputation

**Anonymous Authors**[1]

## Abstract

Automatic Identification System (AIS) trajectories are fundamental to maritime traffic monitoring, vessel behavior analysis, and safety-critical decision support. However, AIS records often contain long missing intervals due to communication failures, limited coverage, harsh maritime environments, or intentional signal disruption. Such long gaps are difficult to reconstruct because sparse boundary observations can admit multiple plausible vessel motions, limiting deterministic interpolation and conventional sequence models. To address this problem, we propose AISFlow, a boundary-informed flow matching framework for long-term vessel trajectory imputation. Unlike diffusion-based approaches that generate trajectories from unstructured Gaussian noise, AISFlow starts from a structured source distribution derived from pre-gap and post-gap vessel states and learns an ODE-based transport toward the missing gap. Boundary-derived kinematic cues, including endpoint displacement, speed/course information, gap length, and heading variation, are further used as conditional information to improve trajectory continuity. AISFlow also supports uncertainty-aware imputation through ensemble sampling and conformal calibration. Experiments on a real-world AIS dataset show that AISFlow achieves state-of-the-art imputation accuracy while reducing inference latency by up to $317\times$ under fixed-step solvers and $9.25\times$ under the adaptive Dopri5 solver, compared with diffusion-based baselines.

## 1. Introduction

Maritime traffic management increasingly relies on vessel trajectory data collected through the Automatic Identification System (AIS). AIS messages provide time-stamped vessel states, including position, speed over ground (SOG), course over ground (COG), and heading, which are essential for maritime surveillance, collision risk assessment, vessel behavior analysis, and safety-critical decision support (Le Tixerant et al., 2018). AIS data streams further enable vessel traffic pattern mining (Yang et al., 2024), and support downstream tasks such as trajectory prediction, anomaly detection, route monitoring, and maritime risk analysis (Nguyen et al., 2018a). Reliable AIS trajectories are therefore a foundational resource across the maritime domain.

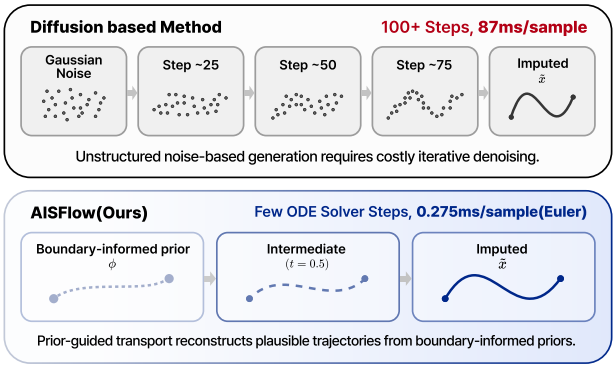


*Figure 1.* Comparison between diffusion-based generation and AISFlow. AISFlow replaces costly iterative denoising from Gaussian noise with faster prior-guided transport from boundary-informed priors. The AISFlow latency shown is based on Euler inference.

However, AIS data are often incomplete in real-world maritime environments. Missing AIS messages can arise from communication failures, limited terrestrial or satellite coverage, equipment malfunction, harsh weather conditions, or intentional signal disruption. While short gaps can often be handled by simple interpolation, long-duration gaps are much more challenging because vessels may change speed, alter heading, wait, maneuver around traffic, or deviate from expected routes. Such gaps can obscure abnormal vessel behaviors, increase uncertainty in collision-risk assessment, and hinder the detection of illicit maritime activities.

Existing AIS imputation approaches can be broadly categorized into deterministic reconstruction, domain-specific

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Submitted to the AI for Science workshop (ICML 2026).

physically constrained methods, and deep learning-based models. Classical methods, such as linear interpolation, spline interpolation, and kinematic reconstruction (Guo et al., 2021), are efficient for locally smooth vessel motion but often fail under long gaps with speed changes, heading variations, or route deviations. Domain-specific methods improve physical plausibility by incorporating maritime motion constraints or vessel kinematics (Sang et al., 2015; Zhang et al., 2018), but they rely on handcrafted assumptions and may not capture complex context-dependent behaviors. Deep learning-based methods further improve AIS imputation by learning movement patterns from historical data, including deterministic sequence models (Cao et al., 2018; Du et al., 2023) and generative approaches (Tashiro et al., 2021; Zhang et al., 2024; Wang et al., 2025). However, diffusion-based generative models require many denoising steps and start from semantically unstructured Gaussian noise, leading to high inference cost and requiring the model to learn a difficult transformation from noise to physically plausible vessel trajectories (see Figure 1).

To address these limitations, we propose **AISFlow**, a **boundary-informed flow matching framework** for long-term vessel trajectory imputation. Instead of generating missing trajectories from unstructured Gaussian noise, AISFlow constructs a structured source distribution from the observed pre-gap and post-gap vessel states. A conditional flow matching model then learns an ODE-based transport from this boundary-informed source to the true missing trajectory, enabling accurate few-step generative imputation. AISFlow further incorporates boundary-derived kinematic cues, including endpoint displacement, SOG/COG-based motion information, and heading variation, as conditional inputs to improve trajectory continuity. We also extend AISFlow with uncertainty-aware imputation through ensemble sampling and conformal calibration, providing calibrated prediction bands for reconstructed AIS trajectories.

The main contributions of this paper are summarized as follows:

- We propose AISFlow, a boundary-informed flow matching framework for long-term AIS trajectory imputation via few-step ODE-based generation.
- We introduce a structured source coupling strategy that initializes generative transport from a boundary-informed source distribution, together with kinematic conditioning to improve trajectory continuity.
- We extend AISFlow with uncertainty-aware imputation through ensemble sampling and conformal calibration, yielding calibrated prediction bands with finite-sample marginal coverage.
- Experiments on real-world AIS data demonstrate state-of-the-art imputation accuracy with up to $317\times$ faster inference than diffusion-based baselines under fixed-step solvers, and $9.25\times$ faster inference under the adaptive Dopri5 solver.

## 2. Related Work

### 2.1. Deterministic and Domain-Specific Methods

Early AIS trajectory reconstruction relies on geometric and kinematic interpolation. Linear and spline interpolation are efficient for locally smooth vessel motion but fail under long gaps with speed changes or heading variations. Gaussian Process-based methods offer a probabilistic extension by modeling positional uncertainty at each timestamp (Nguyen et al., 2023), yet remain limited under long gaps with multi-modal vessel behaviors. Kinematic variables such as SOG, COG, rate of turn (ROT), and change of speed (COS), which are critical for position estimation and collision risk assessment (Sang et al., 2016), form the basis of subsequent reconstruction methods that enforce velocity and acceleration consistency along reconstructed segments (Sang et al., 2015; Zhang et al., 2018; Guo et al., 2021).

More recent methods introduce data-driven elements within deterministic pipelines, including segmentation-based gap-adaptive interpolation (Chen et al., 2025), ensemble interpolation with Transformer-based strategy selection (Li et al., 2025), cleansing-integrated reconstruction under uncertain conditions (Liang et al., 2024), and navigation-state-aware bidirectional kinematic modeling (Liu et al., 2024). Despite these advances, all deterministic methods produce a single output trajectory and rely on handcrafted motion assumptions, providing no calibrated uncertainty estimates for safety-critical maritime decision support.

### 2.2. Deep Sequence Models for Trajectory Imputation

Deep sequence models learn vessel motion patterns from large-scale AIS data and have been applied to various maritime tasks, including trajectory reconstruction, prediction, and anomaly detection (Nguyen et al., 2018a). Sequence-to-sequence models (Nguyen et al., 2018b), recurrent encoder–decoder architectures (Capobianco et al., 2021), and time-series imputation models such as BRITS (Cao et al., 2018) and SAITS (Du et al., 2023) provide strong baselines for recovering missing observations by leveraging temporal context. However, these models are mostly optimized to produce a single point estimate. When long gaps admit multiple plausible trajectories, such objectives may fail to preserve multiple trajectory modes and can bias predictions toward averaged or over-smoothed outputs (Gupta et al., 2018; Amirian et al., 2019). This motivates distributional imputation methods that can represent multiple plausible missing trajectories.

### 2.3. Generative Models for Long-Gap Trajectory Imputation

Generative models address long-gap ambiguity by learning distributions over missing segments rather than producing a single reconstructed trajectory. CSDI (Tashiro et al., 2021) establishes conditional score-based diffusion for probabilistic time-series imputation, and subsequent vessel-specific studies extend diffusion-based imputation to long-gap AIS trajectories through physics-guided diffusion (Zhang et al., 2024) and resampled conditional diffusion (Wang et al., 2025). These methods improve long-gap reconstruction by incorporating historical trajectory embeddings, known trajectory points, or physics-guided losses.

However, diffusion-based imputation requires many denoising steps from Gaussian noise, resulting in high inference cost. Flow matching offers a faster alternative via few-step ODE-based transport (Lipman et al., 2023). AISFlow adapts this paradigm to long-gap AIS imputation by using boundary-derived vessel kinematics as a structured source coupling (Albergo et al., 2024). We further complement point-wise reconstruction accuracy with calibrated trajectory-level uncertainty estimation, which remains less explored in vessel trajectory imputation.

## 3. Problem Formulation

Trajectory imputation aims to recover missing trajectories from partially observed vessel trajectories. We consider an AIS trajectory as a sequence $\mathbf{X} = \{\mathbf{s}_1, \ldots, \mathbf{s}_T\}$, where each state $\mathbf{s}_t = (\mathrm{lat}_t, \mathrm{lon}_t, \mathrm{sog}_t, \mathrm{cog}_t)$ contains the vessel's geographic coordinates, SOG, and COG at timestamp $t$. Given a continuous long missing gap of length $L$ starting at timestamp $k$, a trajectory sequence $\mathbf{X}$ can be decomposed into the observed trajectory sequence $\mathbf{x}_{obs} = [\mathbf{s}_1, \ldots, \mathbf{s}_{k-1}, \mathbf{s}_{k+L}, \ldots, \mathbf{s}_T] \in \mathbb{R}^{L\times 4}$ and the corresponding ground-truth missing sequence as $\mathbf{x}_1 = [\mathbf{s}_k, \ldots, \mathbf{s}_{k+L-1}]$.

The model observes only the boundary states surrounding the missing gap, summarized as a conditioning variable $\mathbf{c}$. This context includes endpoint coordinates, boundary kinematic attributes, and the gap duration.

$$q_{data}(\mathbf{x}_1) \approx p_\theta(\mathbf{x}_1 \mid \mathbf{c})$$

Our goal is to learn the conditional distribution $p_\theta(\mathbf{x}_1 \mid \mathbf{c})$ that approximates the true distribution of missing AIS trajectories $q_{data}(\mathbf{x}_1)$. While the reconstructed gap should remain consistent with the observed boundary information, the same boundary states may still admit multiple plausible vessel trajectories. Therefore, We formulate long-term AIS trajectory imputation as probabilistic generation problem rather than deterministic interpolation.

## 4. Methodology

AISFlow is a boundary-informed flow-matching framework for long-term AIS trajectory imputation. Given the boundary context $\mathbf{c}$, AISFlow constructs a trajectory-valued source from the observed boundaries and learns a transport path from the informed source $\mathbf{x}_0$ to the ground-truth missing AIS trajectory $\mathbf{x}_1 \in \mathbb{R}^{L\times 4}$.

The framework consists of four components: boundary-informed source coupling, conditional flow matching, a context-conditioned vector field network, and calibrated uncertainty estimation. Unlike diffusion-based imputation methods that generate trajectories through many denoising steps from Gaussian noise (Zhang et al., 2024; Wang et al., 2025), AISFlow starts from a boundary-consistent source trajectory and refines it through few-step ODE sampling (Lipman et al., 2023).

### 4.1. Boundary-Informed Source Coupling

A key in flow matching is the source distribution from which generative transport starts. Flow-matching for generative modeling often uses an independent Gaussian source, which is easy to sample but carries no trajectory-specific structure. For long-term AIS imputation, however, the observed boundary states provide informative constraints on the missing gap. AISFlow therefore replaces the independent Gaussian source with a boundary-informed source. This perspective is related to recent work on data-dependent couplings in stochastic interpolants (Albergo et al., 2024).

Given boundary context $\mathbf{c}$, AISFlow defines a deterministic source trajectory $\phi(\mathbf{c}) \in \mathbb{R}^{L\times 4}$, and constructs the source trajectory as

$$\mathbf{x}_0 = \phi(\mathbf{c}) + \sigma\boldsymbol{\epsilon}, \qquad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}).$$

This yields

$$p_0(\mathbf{x}_0 \mid \mathbf{c}) = \mathcal{N}\big(\phi(\mathbf{c}), \sigma^2\mathbf{I}\big)\,.$$

During training, the target trajectory provides $\mathbf{x}_1$, while $\mathbf{x}_0$ is sampled from $p_0(\mathbf{x}_0 \mid \mathbf{c})$. Thus,

$$\pi(\mathbf{x}_0, \mathbf{x}_1) = p_1(\mathbf{x}_1)p_0(\mathbf{x}_0 \mid \mathbf{c}).$$

$\pi(\mathbf{x}_0, \mathbf{x}_1)$ denotes a coupling distribution between the source distribution $p_0$ and the target data distribution $p_1$. This makes the source distribution adapted to the observed boundary states rather than sampled from a data-agnostic Gaussian. The perturbation scale $\sigma$ controls the spread around the source trajectory, acting as a stochastic regularizer and enabling ensemble-based uncertainty estimation. A simple target-energy decomposition supporting this intuition is given in Appendix A.

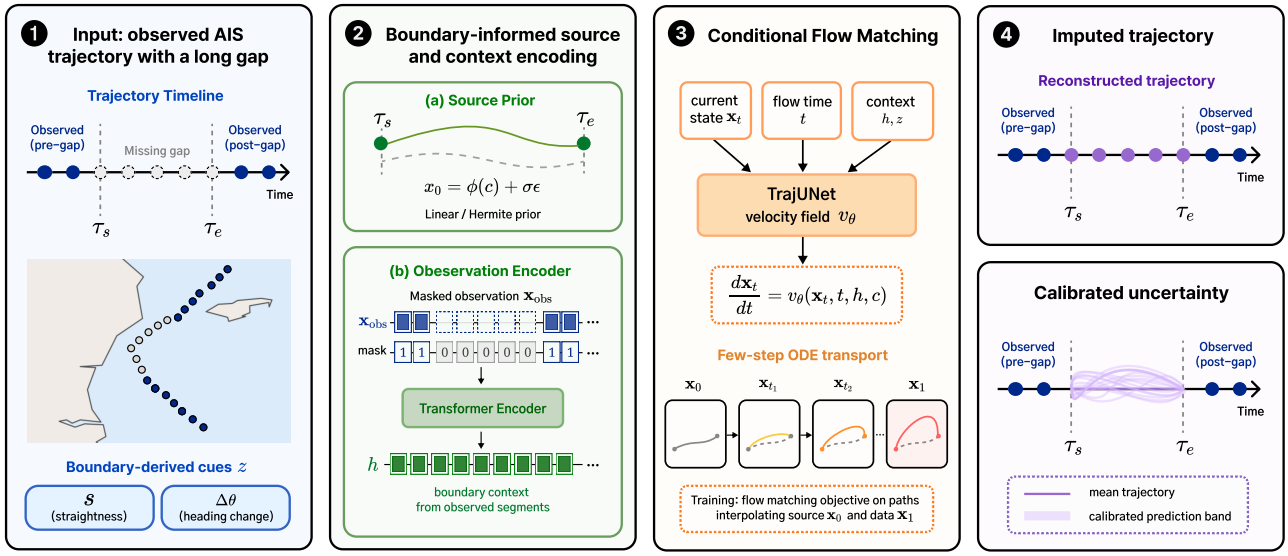


*Figure 2.* AISFlow pipeline. The model uses pre- and post-gap AIS observations to construct a boundary-informed source, encodes observed context, and performs conditional flow matching to reconstruct the missing segment with calibrated uncertainty.

We instantiate $\phi(\mathbf{c})$ using boundary-consistent interpolation sources. The linear prior bridges the missing gap by interpolation and serves as a simple trajectory source. The Hermite prior additionally incorporates boundary kinematics. For the position channels, let $\mathbf{p}_s$ and $\mathbf{p}_e$ denote the start and end boundary positions, and let $\mathbf{v}_s$ and $\mathbf{v}_e$ denote boundary velocity vectors converted from SOG and COG and scaled to the normalized gap time. For normalized gap time $\tau \in [0, 1]$, the cubic Hermite position prior is

$$p_{\text{pos}}(\tau) = h_{00}(\tau)\mathbf{p}_s + h_{10}(\tau)\mathbf{v}_s + h_{01}(\tau)\mathbf{p}_e + h_{11}(\tau)\mathbf{v}_e,$$

where $h_{00}, h_{10}, h_{01}, h_{11}$ are the standard cubic Hermite basis functions. The SOG channel is interpolated linearly, and the COG channel is interpolated using circular interpolation to respect angular periodicity.

This construction initializes the ODE from a trajectory-shaped source, allowing flow matching to focus on refining the boundary prior toward the missing AIS segment.

### 4.2. Conditional Flow Matching and ODE Sampling

Given a target missing AIS sequence $\mathbf{x}_1$ and a boundary-informed source trajectory $\mathbf{x}_0$, AISFlow learns a conditional velocity field that transports $\mathbf{x}_0$ toward $\mathbf{x}_1$. We use a linear interpolation path between the source and target:

$$\mathbf{x}_t = (1 - t)\mathbf{x}_0 + t\mathbf{x}_1, \qquad t \sim \mathcal{U}(0, 1).$$

The corresponding target velocity is

$$\mathbf{u} = \mathbf{x}_1 - \mathbf{x}_0.$$

The conditional flow-matching objective is

$$\mathcal{L}_{\text{FM}} = \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_1) \sim \pi,\, t \sim \mathcal{U}(0,1)} \left[ \| v_\theta(\mathbf{x}_t, t, \mathbf{c}) - \mathbf{u} \|_2^2 \right],$$

where $\pi$ is the training joint distribution defined in Section 4.1.

At inference, AISFlow generates an imputed trajectory by solving the learned ODE

$$\frac{d\mathbf{x}_t}{dt} = v_\theta(\mathbf{x}_t, t, \mathbf{c}), \qquad \mathbf{x}_{t=0} = \mathbf{x}_0.$$

For deterministic point prediction, we initialize the ODE from the deterministic boundary prior, i.e., $\mathbf{x}_0 = \phi(\mathbf{c})$. For uncertainty estimation, we sample multiple perturbed sources from $p_0(\mathbf{x}_0 \mid \mathbf{c})$ and solve the ODE independently. Although the formulation above focuses on the missing segment, our implementation operates on the full trajectory window and applies the loss and ODE updates only to the masked gap region.

### 4.3. Network Architecture and Boundary Conditioning

AISFlow uses two neural components: an observation encoder and a velocity network. Given a partially observed AIS trajectory and its binary observation mask, the observation encoder produces a contextual feature sequence $\mathbf{h} \in \mathbb{R}^{T \times d_h}$ aligned with the trajectory timesteps. It embeds the position and motion channels using learnable Fourier features (Li et al., 2021; Zhang et al., 2024), adds temporal positional encodings, and applies a Transformer encoder with masked attention. The attention mask prevents missing tokens from contributing to the context encoding. The resulting representation provides the velocity network with contextual information from the observed pre-gap and post-gap trajectory segments.

The velocity field $v_\theta$ is parameterized by a one-dimensional

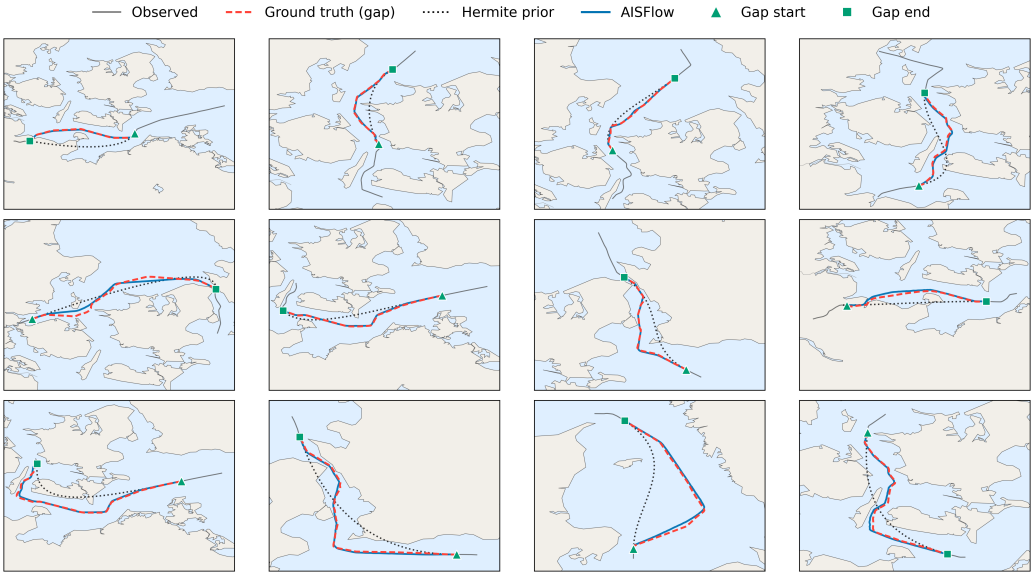


*Figure 3.* Qualitative examples of long-gap AIS trajectory imputation. Rows correspond to 6h, 8h, and 10h gap cases from top to bottom. Each panel shows the observed pre- and post-gap trajectory segments, the ground-truth missing gap, the Hermite prior, and the AISFlow prediction. Across diverse vessel motions and geographic regions, AISFlow closely follows the ground-truth gap trajectory while preserving consistency with the observed boundary points.

U-Net over the trajectory time axis. At each flow time $t$, the network receives the current trajectory state $\mathbf{x}_t$, the encoded observation context $\mathbf{h}$, the gap mask, and trajectory-level boundary cues, and outputs a velocity with the same dimension as the AIS state sequence. The flow-time embedding is injected into each residual block of the U-Net. Scalar boundary cues are injected through a small MLP. We define a boundary straightness cue

$$s = \frac{\|\mathbf{p}_e - \mathbf{p}_s\|}{D_{\text{exp}}},$$

where $\|\mathbf{p}_e - \mathbf{p}_s\|$ denotes the displacement between the observed gap endpoints, and $D_{\text{exp}}$ is the expected travel distance estimated by multiplying the average boundary SOG by the gap duration. We also use the wrapped boundary heading change

$$\Delta\theta = \text{COG}_e - \text{COG}_s.$$

Together, $s$ and $\Delta\theta$ indicate whether the missing interval is likely to be nearly straight or involve substantial turning, allowing the velocity field to adapt its refinement behavior.

This separation allows AISFlow to encode boundary observations once and repeatedly evaluate the U-Net during ODE sampling to refine the boundary-informed source trajectory.

### 4.4. Calibrated Uncertainty Estimation

AISFlow estimates uncertainty by sampling multiple boundary-informed source trajectories. Given the same boundary context $\mathbf{c}$, we draw $N$ independent perturbations

$$\mathbf{x}_0^{(k)} = \phi(\mathbf{c}) + \sigma\boldsymbol{\epsilon}^{(k)}, \quad \boldsymbol{\epsilon}^{(k)} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}),\ k = 1, \ldots, N,$$

and solve the ODE for each source to obtain an ensemble of imputations $\{\hat{\mathbf{x}}_1^{(k)}\}_{k=1}^N$. We summarize the ensemble using its element-wise mean and standard deviation:

$$\hat{\boldsymbol{\mu}} = \frac{1}{N}\sum_{k=1}^{N} \hat{\mathbf{x}}_1^{(k)}, \quad \hat{\boldsymbol{\sigma}} = \sqrt{\frac{1}{N-1}\sum_{k=1}^{N}\left(\hat{\mathbf{x}}_1^{(k)} - \hat{\boldsymbol{\mu}}\right)^2}.$$

Since the raw ensemble spread is not necessarily calibrated, we apply split conformal prediction on a held-out calibration set (Angelopoulos & Bates, 2023; Stankeviciute et al., 2021). For each calibration trajectory, timestep, and spatial channel, we compute the normalized residual

$$z_i^{t,j} = \frac{\left|\mathbf{x}_{1,i}^{t,j} - \hat{\boldsymbol{\mu}}_i^{t,j}\right|}{\hat{\boldsymbol{\sigma}}_i^{t,j} + \eta},$$

where $t$ indexes missing timesteps, $j$ indexes calibrated spatial channels, and $\eta$ is a small constant for numerical stability. We pool the element-wise residuals $\{z_i^{t,j}\}$ across calibration trajectories, timesteps, and channels, and compute the conformal quantile $\hat{q}_{1-\alpha}$.

For a test trajectory, the calibrated prediction interval is

$$\mathcal{I}^{t,j} = \left[\hat{\boldsymbol{\mu}}^{t,j} - \hat{q}_{1-\alpha}\hat{\boldsymbol{\sigma}}^{t,j},\ \hat{\boldsymbol{\mu}}^{t,j} + \hat{q}_{1-\alpha}\hat{\boldsymbol{\sigma}}^{t,j}\right].$$

This marginal calibration rescales AISFlow's ensemble spread into compact prediction intervals with distribution-free marginal coverage over missing timesteps and spatial channels.

# 5. Experiments

## 5.1. Dataset

We evaluate AISFlow on a publicly available AIS dataset provided by the Danish Maritime Authority.[1] We extract AIS observations of cargo and tanker vessels recorded from July 2022 to March 2023 within a geographic region spanning $53.5°$–$59.0°$N latitude and $7.5°$–$14.0°$E longitude. During preprocessing, we remove abnormal speed records ($\geq 30$ knots) and stationary vessels (moored or at anchor). Non-contiguous voyages are split into contiguous segments using a maximum inter-message interval of 2 hours, and trajectories shorter than 20 points or 4 hours in duration are discarded. Each trajectory is downsampled to 5-minute intervals and segmented into 12-hour windows, yielding a fixed sequence length of 144. We partition the dataset into training, validation, and test sets with a ratio of 7:1:2.

## 5.2. Evaluation Protocol and Experimental Setup

**Evaluation Metrics.** We evaluate imputation performance using Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) computed separately for latitude (LAT) and longitude (LON) of reconstructed vessel positions.

**Gap Simulation.** To simulate long-term AIS signal loss, we apply a continuous masking strategy over each trajectory sample. The mask length is varied across $\{6\text{h}, 8\text{h}, 10\text{h}\}$ to assess model robustness under increasingly challenging imputation scenarios. The start position of each mask is sampled uniformly at random, and all trajectory observations within the masked interval are treated as completely missing during both training and evaluation.

**Experimental Setup.** AISFlow is trained for 200 epochs using Adam with a learning rate of $2 \times 10^{-4}$ and a batch size of 256. Unless otherwise specified, we use the Hermite boundary prior with $\sigma = 0.10$ and Dopri5 for ODE-based inference. For uncertainty estimation, we use $N = 16$ source perturbations. All baselines are evaluated using the same data splits and gap simulation protocol. Full architecture, implementation, and hardware details are provided in Appendix B.

## 5.3. Baselines

We compare AISFlow against the following baselines:

- **LI**: Linear interpolation between observed boundary points.
- **TraImpute** (Elshrif et al., 2022): A network-free heuristic search algorithm that infers missing locations based on frequency patterns observed in historical trajectories.
- **Transformer** (Vaswani et al., 2017): A seq2seq model with self-attention for autoregressive trajectory decoding.
- **Multi-task AIS** (Nguyen et al., 2018a): A recurrent network with latent variable modeling for joint AIS representation learning and vessel trajectory reconstruction.
- **CSDI** (Tashiro et al., 2021): A conditional score-based diffusion model for probabilistic time series imputation conditioned on observed data.
- **LT** (Zhang et al., 2024): A physics-guided diffusion model for long-term vessel trajectory imputation that incorporates kinematic constraints and historical trajectory embeddings as conditional information.
- **TrajDiff** (Wang et al., 2025): A cyclic denoising diffusion model that integrates known trajectory points via a conditional embedding mechanism to guide missing segment generation.

## 5.4. Results

Table 1 compares AISFlow with existing vessel trajectory imputation methods under all missing scenarios(6h, 8h, 10h). Table 2 evaluates inference efficiency between LT and AISFlow under different ODE solvers. Tables 3 and 4 summarize ablation results on source prior selection and kinematic conditioning, respectively. Additional sensitivity results for the source perturbation scale $\sigma$ are provided in Appendix A.2.

**Imputation Performance (Table 1).** AISFlow attains the lowest MAE and RMSE across all missing-gap scenarios, with a substantial margin over the best-performing baseline. These results are consistent with the motivation of AISFlow: coupling flow matching with a boundary-informed prior tailored to vessel trajectory imputation. Rather than transporting samples from a Gaussian prior, AISFlow initializes the flow from a prior that already encodes the observed boundary context of each missing gap. This prior coupling provides trajectory-specific context and allows the velocity field to model only the residual deviation from a kinematically plausible interpolant. Consequently, AISFlow is encouraged to predict trajectories that are both accurate and

[1] https://dma.dk/safety-at-sea/navigational-information/ais-data

*Table 1.* Trajectory imputation performance across gap lengths (6h, 8h, 10h). Each cell reports MAE / RMSE for LAT and LON. **Bold** and underline denote the best and second-best results.

| Model | 6h | | 8h | | 10h | |
|---|---|---|---|---|---|---|
| | LAT | LON | LAT | LON | LAT | LON |
| LI | 0.0612 / 0.0902 | 0.0997 / 0.1444 | 0.0920 / 0.1323 | 0.1406 / 0.1989 | 0.1246 / 0.1765 | 0.1859 / 0.2608 |
| TraImpute | 0.0439 / 0.0843 | 0.0761 / 0.1289 | 0.0538 / 0.1233 | 0.0851 / 0.1916 | 0.0629 / 0.1242 | 0.1051 / 0.1681 |
| Transformer | 0.0433 / 0.0835 | 0.0749 / 0.1089 | 0.0475 / 0.0901 | 0.0909 / 0.1609 | 0.0518 / 0.0824 | 0.0893 / 0.1457 |
| Multi-task AIS | 0.0427 / 0.0810 | 0.0773 / 0.1197 | 0.0436 / 0.0833 | 0.0821 / 0.1709 | 0.0642 / 0.0944 | 0.1003 / 0.1566 |
| CSDI | 0.0328 / 0.0468 | 0.0542 / 0.0819 | 0.0391 / 0.0572 | 0.0755 / 0.1223 | 0.0492 / 0.0862 | 0.0832 / 0.1311 |
| LT-NR | 0.0307 / 0.0459 | 0.0504 / 0.0796 | 0.0369 / 0.0516 | 0.0639 / 0.0987 | 0.0420 / 0.0671 | 0.0793 / 0.1243 |
| LT-NP | 0.0268 / 0.0390 | 0.0482 / 0.0725 | 0.0329 / 0.0506 | 0.0606 / 0.0974 | 0.0374 / 0.0602 | 0.0625 / 0.1050 |
| LT | 0.0296 / 0.0424 | 0.0478 / 0.0723 | 0.0324 / 0.0500 | 0.0592 / 0.0962 | 0.0366 / 0.0593 | 0.0607 / 0.1036 |
| TrajDiff_C | 0.0170 / 0.0250 | 0.0310 / 0.0440 | 0.0230 / 0.0320 | 0.0410 / 0.0850 | 0.0280 / 0.0380 | 0.0690 / 0.1040 |
| TrajDiff_R | 0.0520 / 0.0660 | 0.0700 / 0.0910 | 0.0640 / 0.0830 | 0.0760 / 0.1150 | 0.0900 / 0.1140 | 0.1480 / 0.2050 |
| TrajDiff | 0.0160 / 0.0230 | 0.0270 / 0.0410 | 0.0200 / 0.0280 | 0.0400 / 0.0560 | 0.0260 / 0.0360 | 0.0610 / 0.0890 |
| **AISFlow(Ours)** | **0.0087 / 0.0171** | **0.0157 / 0.0333** | **0.0120 / 0.0232** | **0.0215 / 0.0438** | **0.0165 / 0.0315** | **0.0301 / 0.0615** |

consistent with the motion implied by the observed boundaries. This makes the learning problem easier and leads to stronger imputation performance.

*Table 2.* Inference efficiency and imputation performance(6h) under different ODE solvers.

| Method | Latency(ms) | Speedup | LAT | LON |
|---|---|---|---|---|
| LT | 87.2 | 1.0× | 0.0296 / 0.0424 | 0.0478 / 0.0723 |
| Ours (Euler) | 0.275 | 317× | 0.0101 / 0.0179 | 0.0174 / 0.0339 |
| Ours (Midpoint) | 0.436 | 200× | 0.0095 / 0.0175 | 0.0167 / 0.0337 |
| Ours (RK4) | 0.750 | 116.27× | 0.0215 / 0.0293 | 0.0262 / 0.0408 |
| **Ours (Dopri5)** | **9.428** | **9.25×** | **0.0087 / 0.0171** | **0.0157 / 0.0333** |

**Inference Efficiency (Table 2).** Inference efficiency is measured using inference latency averaged over all test samples, and speedup is reported relative to LT, which requires 87.2 ms per sample. The number of ODE solver steps is fixed to 3 for fixed-step solvers, whereas Dopri5 uses adaptive step sizes. As shown in Table 2, all AISFlow variants substantially outperform LT in both latency and imputation accuracy. Euler achieves the best inference efficiency with a 317× speedup over LT, while Dopri5 achieves the strongest imputation performance. This reflects a natural solver trade-off: Euler has minimal computational overhead and is suitable for latency-sensitive settings, while Dopri5 more accurately integrates the learned velocity field through adaptive higher-order updates. In this work, we adopt Dopri5 as the default solver for AISFlow to prioritize accurate trajectory prediction under long missing gaps.

**Source Prior Ablation (Table 3).** Structured boundary priors substantially outperform the unstructured Gaussian source across all gap lengths, confirming that boundary-informed initialization reduces the difficulty of generative transport. Among structured priors, Linear and Cubic Hermite yield statistically indistinguishable performance (gap-averaged MAE difference $\leq 0.6\%$, $p > 0.49$, paired $t$-test),

*Table 3.* Prior family ablation ($\sigma$=0.10, Dopri5, 3-seeds average). Each cell reports MAE / RMSE averaged over LAT and LON. **Bold**: best, underline: second-best.

| Prior | 6h | 8h | 10h |
|---|---|---|---|
| Linear | **0.0087/0.0249** | **0.0119/0.0334** | **0.0163**/0.0467 |
| Cubic Hermite† | **0.0087**/0.0252 | 0.0120/0.0335 | 0.0165/**0.0465** |
| CTRA | 0.0102/0.0275 | 0.0141/0.0378 | 0.0193/0.0531 |
| Gaussian | 0.0151/0.0372 | 0.0199/0.0480 | 0.0264/0.0648 |
| Quintic Hermite | 0.0371/0.0761 | 0.0457/0.0943 | 0.0584/0.1187 |

† Selected as default prior.

likely because 5-minute AIS sampling provides sufficiently dense boundary observations for both to produce similar initializations. Higher-order priors (CTRA, Quintic Hermite) degrade sharply, as finite-difference estimation of acceleration and curvature amplifies noise at the $1/\Delta t^2$ scale. We adopt Cubic Hermite as the default prior, as it explicitly encodes boundary velocity via observed SOG/COG and is expected to be more discriminative in sparser sampling.

*Table 4.* Boundary conditioning ablation (Dopri5, 3-seed average). Each cell reports MAE / RMSE. $s$: endpoint displacement; $\Delta\theta$: heading variation.

| | | None | $s$ only | $\Delta\theta$ only | $s$+$\Delta\theta$ |
|---|---|---|---|---|---|
| **6h** | LAT | 0.0088/0.0170 | **0.0087**/0.0170 | 0.0088/0.0171 | **0.0087/0.0168** |
| | LON | 0.0159/0.0333 | **0.0156**/0.0329 | 0.0158/0.0332 | **0.0156/0.0326** |
| **8h** | LAT | 0.0121/0.0232 | **0.0119**/0.0230 | 0.0120/0.0231 | **0.0119/0.0229** |
| | LON | 0.0216/0.0430 | 0.0213/0.0429 | 0.0215/0.0431 | **0.0212/0.0427** |
| **10h** | LAT | 0.0165/0.0312 | 0.0163/0.0311 | 0.0164/0.0311 | **0.0162/0.0310** |
| | LON | 0.0302/0.0605 | **0.0298/0.0604** | 0.0300/0.0606 | **0.0298**/0.0607 |

**Boundary Conditioning Ablation (Table 4).** Adding kinematic boundary cues consistently improves imputation over the unconditioned baseline. The full conditioning ($s$+$\Delta\theta$) achieves the best performance in most settings, with notable gains in LON at 6h and 8h. At 10h LON, $s$

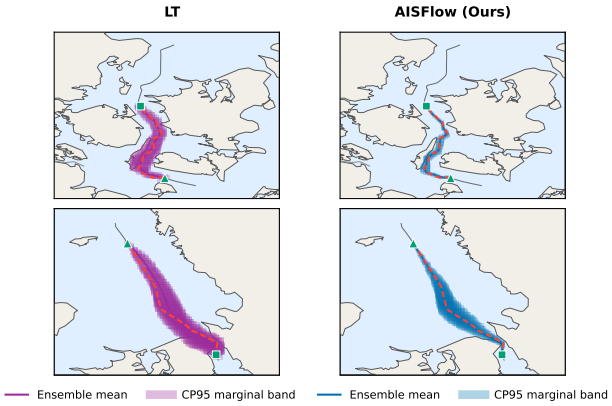


*Figure 4.* Qualitative comparison of calibrated CP95 marginal bands. AISFlow produces tighter uncertainty bands than the diffusion baseline while covering plausible vessel trajectories.

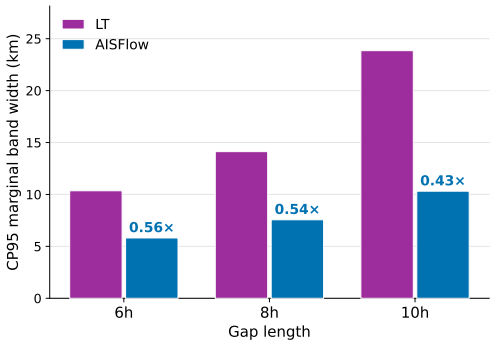


*Figure 5.* Mean CP95 marginal band width across gap lengths. AISFlow requires substantially narrower calibrated bands than the diffusion baseline.

alone performs on par or better, suggesting that displacement information is the primary driver for longer gaps. The absolute improvement is marginal on average ($\sim$1%), but per-sample analysis shows 13–25% gains for moderate-turn trajectories ($|\Delta\theta| \in [60°, 90°]$). These results indicate that boundary-derived kinematic cues effectively leverage the unmasked pre-gap and post-gap context, enabling more consistent trajectory reconstruction across the missing interval.

### 5.5. Calibrated Uncertainty Quantification

We evaluate AISFlow's calibrated uncertainty estimates using marginal conformal prediction with a nominal coverage level of $1 - \alpha = 0.95$, denoted as marginal CP95. AISFlow generates trajectory ensembles by sampling stochastic perturbations around the boundary-informed source, and split conformal prediction rescales the ensemble spread into calibrated uncertainty bands.

Fig. 4 compares marginal CP95 bands from AISFlow and the diffusion baseline. Both methods are calibrated to the same nominal level, but AISFlow produces substantially tighter bands that remain concentrated around plausible vessel paths. The diffusion baseline requires wider bands, especially around turning regions and long missing intervals. Fig. 5 summarizes the mean marginal CP95 band width across different gap lengths. AISFlow consistently achieves narrower calibrated bands, using only $0.56\times$, $0.54\times$, and $0.43\times$ of the baseline width for 6h, 8h, and 10h gaps, respectively. This indicates that AISFlow provides sharper uncertainty estimates while maintaining calibrated marginal coverage.

## 6. Conclusion

We presented AISFlow, a boundary-informed flow matching framework for long-term AIS trajectory imputation. By initializing generative transport from a boundary-consistent source trajectory rather than unstructured Gaussian noise, AISFlow reduces the difficulty of the learning problem and enables accurate few-step ODE-based generation. Boundary-derived kinematic cues further improve trajectory continuity, and conformal calibration provides coverage-guaranteed uncertainty bands. Experiments on a real-world AIS dataset demonstrate state-of-the-art imputation accuracy alongside substantially reduced inference latency. AISFlow is up to $317\times$ faster than diffusion-based baselines under fixed-step solvers and $9.25\times$ faster under the adaptive Dopri5 solver. This efficiency makes AISFlow a practical solution for real-time maritime traffic management, where timely trajectory reconstruction is essential for collision risk assessment and situational awareness.

## 7. Limitations and Future Work

AISFlow has several limitations. First, evaluation is conducted on a single regional AIS dataset, and generalization to other maritime zones, vessel types, or sampling rates remains to be validated. Second, external contextual factors such as weather conditions, sea state, and traffic density are not incorporated, which may limit imputation plausibility in complex scenarios. Third, conformal calibration does not model spatial correlation structure of prediction errors across timesteps. Future work includes multi-region benchmarking, integration of environmental context, and exploration of structured uncertainty representations for safety-critical applications.

## References

Albergo, M. S., Goldstein, M., Boffi, N. M., Ranganath, R., and Vanden-Eijnden, E. Stochastic interpolants with data-dependent couplings. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.

Amirian, J., Hayet, J.-B., and Pettré, J. Social ways: Learn-

ing multi-modal distributions of pedestrian trajectories with gans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 0–0, 2019.

Angelopoulos, A. N. and Bates, S. Conformal prediction: A gentle introduction. *Foundations and Trends in Machine Learning*, 16(4):494–591, 2023. doi: 10.1561/2200000101.

Cao, W., Wang, D., Li, J., Zhou, H., Li, L., and Li, Y. Brits: Bidirectional recurrent imputation for time series. *Advances in neural information processing systems*, 31, 2018.

Capobianco, S., Millefiori, L. M., Forti, N., Braca, P., and Willett, P. Deep learning methods for vessel trajectory prediction based on recurrent neural networks. *IEEE Transactions on Aerospace and Electronic Systems*, 57 (6):4329–4346, 2021.

Chen, Y., Chen, Y., Cui, Y., Cai, X., Yin, C., and Cheng, Y. Optimizing vessel trajectories: Advanced denoising and interpolation techniques for ais data. *Ocean Engineering*, 327:120988, 2025.

Du, W., Côté, D., and Liu, Y. Saits: Self-attention-based imputation for time series. *Expert Systems with Applications*, 219:119619, 2023.

Elshrif, M. M., Isufaj, K., and Mokbel, M. F. Networkless trajectory imputation. In *Proceedings of the 30th International Conference on Advances in Geographic Information Systems*, pp. 1–10, 2022.

Guo, S., Mou, J., Chen, L., and Chen, P. Improved kinematic interpolation for ais trajectory reconstruction. *Ocean Engineering*, 234:109256, 2021.

Gupta, A., Johnson, J., Fei-Fei, L., Savarese, S., and Alahi, A. Social gan: Socially acceptable trajectories with generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

Le Tixerant, M., Le Guyader, D., Gourmelon, F., and Queffelec, B. How can automatic identification system (ais) data be used for maritime spatial planning? *Ocean & Coastal Management*, 166:18–30, 2018.

Li, Y., Si, S., Li, G., Hsieh, C.-J., and Bengio, S. Learnable fourier features for multi-dimensional spatial positional encoding. *Advances in Neural Information Processing Systems*, 34:15816–15829, 2021.

Li, Y., Guo, S., Chen, P., Chen, L., and Mou, J. A stacking-based ensemble learning model for intelligent ship trajectory interpolation. *Reliability Engineering & System Safety*, pp. 111615, 2025.

Liang, M., Su, J., Liu, R. W., and Lam, J. S. L. Aisclean: Ais data-driven vessel trajectory reconstruction under uncertain conditions. *Ocean Engineering*, 306:117987, 2024.

Lipman, Y., Chen, R. T. Q., Ben-Hamu, H., Nickel, M., and Le, M. Flow matching for generative modeling. In *International Conference on Learning Representations*, 2023.

Liu, Y., Shi, Z., Fu, B., Xu, H., and Wu, H. Marine trajectory reconstruction method based on navigation state recognition and bi-directional kinematic interpolation. *Journal of Marine Science and Engineering*, 12(12):2164, 2024.

Nguyen, D., Vadaine, R., Hajduch, G., Garello, R., and Fablet, R. A multi-task deep learning architecture for maritime surveillance using ais data streams. In *2018 IEEE 5th International Conference on Data Science and Advanced Analytics (DSAA)*, pp. 331–340. IEEE, 2018a.

Nguyen, D.-D., Le Van, C., and Ali, M. I. Vessel trajectory prediction using sequence-to-sequence models over spatial grid. In *Proceedings of the 12th ACM international conference on distributed and event-based systems*, pp. 258–261, 2018b.

Nguyen, K., Krumm, J., and Shahabi, C. Gaussian process for trajectories. In *Spatial Gems, Volume 2*, pp. 37–48. 2023.

Sang, L.-z., Wall, A., Mao, Z., Yan, X.-p., and Wang, J. A novel method for restoring the trajectory of the inland waterway ship by using ais data. *Ocean Engineering*, 110: 183–194, 2015.

Sang, L.-z., Yan, X.-p., Wall, A., Wang, J., and Mao, Z. Cpa calculation method based on ais position prediction. *The journal of navigation*, 69(6):1409–1426, 2016.

Stankeviciute, K., Alaa, A. M., and van der Schaar, M. Conformal time-series forecasting. In *Advances in Neural Information Processing Systems*, volume 34, pp. 6216–6228, 2021.

Tashiro, Y., Song, J., Song, Y., and Ermon, S. Csdi: Conditional score-based diffusion models for probabilistic time series imputation. *Advances in neural information processing systems*, 34:24804–24816, 2021.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

Wang, W., Xiong, W., Chen, L., and Chen, H. Trajdiff: A method for vessel trajectory imputation utilizing resampled conditional diffusion models. *Results in Engineering*, pp. 107279, 2025.

Yang, J., Bian, X., Qi, Y., Wang, X., Yang, Z., and Liu, J. A spatial-temporal data mining method for the extraction of vessel traffic patterns using ais data. *Ocean Engineering*, 293:116454, 2024.

Zhang, L., Meng, Q., Xiao, Z., and Fu, X. A novel ship trajectory reconstruction approach using ais data. *Ocean Engineering*, 159:165–174, 2018.

Zhang, Z., Fan, Z., Lv, Z., Song, X., and Shibasaki, R. Long-term vessel trajectory imputation with physics-guided diffusion probabilistic model. In *Proceedings of the 30th ACM SIGKDD conference on knowledge discovery and data mining*, pp. 4398–4407, 2024.

# A. Coupling-Perturbation Analysis

This section provides a simple target-energy analysis for the boundary-informed source coupling used in AISFlow. The analysis does not directly guarantee lower final imputation error, but it shows that a structured source prior can reduce the average displacement that the flow-matching velocity field must learn.

## A.1. Target-Energy Decomposition

Consider a missing trajectory segment vectorized as $\mathbf{x}_1 \in \mathbb{R}^d$, where $d$ denotes the number of reconstructed coordinates in the masked region. Given boundary context $\mathbf{c}$, AISFlow constructs a deterministic boundary-informed prior $\phi(\mathbf{c}) \in \mathbb{R}^d$ and samples the perturbed source as

$$\mathbf{x}_0^{(\phi)} = \phi(\mathbf{c}) + \sigma\boldsymbol{\epsilon}, \qquad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d),$$

where $\boldsymbol{\epsilon}$ is independent of $(\mathbf{c}, \mathbf{x}_1)$. For the linear flow-matching path, the target velocity is

$$\mathbf{u}^{(\phi)} = \mathbf{x}_1 - \mathbf{x}_0^{(\phi)} = \mathbf{x}_1 - \phi(\mathbf{c}) - \sigma\boldsymbol{\epsilon}.$$

**Proposition A.1.** *For the perturbed boundary-informed source above, the expected target energy satisfies*

$$\mathbb{E}\left[\left\|\mathbf{u}^{(\phi)}\right\|_2^2\right] = \mathbb{E}\left[\|\mathbf{x}_1 - \phi(\mathbf{c})\|_2^2\right] + \sigma^2 d.$$

*Proof.* Let $\mathbf{a} = \mathbf{x}_1 - \phi(\mathbf{c})$. Conditioned on $(\mathbf{c}, \mathbf{x}_1)$, $\mathbf{a}$ is fixed and

$$\left\|\mathbf{u}^{(\phi)}\right\|_2^2 = \|\mathbf{a} - \sigma\boldsymbol{\epsilon}\|_2^2 = \|\mathbf{a}\|_2^2 - 2\sigma\mathbf{a}^\top\boldsymbol{\epsilon} + \sigma^2\|\boldsymbol{\epsilon}\|_2^2.$$

Taking expectation over $\boldsymbol{\epsilon}$ and using $\mathbb{E}[\boldsymbol{\epsilon}] = \mathbf{0}$ and $\mathbb{E}\|\boldsymbol{\epsilon}\|_2^2 = d$, we obtain

$$\mathbb{E}_{\boldsymbol{\epsilon}}\left[\left\|\mathbf{u}^{(\phi)}\right\|_2^2 \mid \mathbf{c}, \mathbf{x}_1\right] = \|\mathbf{x}_1 - \phi(\mathbf{c})\|_2^2 + \sigma^2 d.$$

Finally, taking expectation over $(\mathbf{c}, \mathbf{x}_1) \sim p_{\mathrm{data}}$ gives

$$\mathbb{E}\left[\left\|\mathbf{u}^{(\phi)}\right\|_2^2\right] = \mathbb{E}\left[\|\mathbf{x}_1 - \phi(\mathbf{c})\|_2^2\right] + \sigma^2 d.$$

□

The result shows that the flow-matching target energy decomposes into a deterministic prior error term and an injected noise energy term. When the boundary-informed prior $\phi(\mathbf{c})$ is close to the true missing segment $\mathbf{x}_1$, the deterministic term decreases, and the velocity field learns a shorter average transport displacement. In contrast, an unstructured Gaussian source can be viewed as the special case where the source center does not depend on the boundary context, so its target energy is determined by the distance from noise to the full trajectory rather than by the boundary-to-target prior error.

The decomposition should be interpreted as a source-design rationale rather than a direct guarantee of lower MAE. It shows that boundary-informed coupling can reduce the expected transport burden when the prior is close to the missing trajectory. This supports the source-coupling design of AISFlow: boundary-informed sources reduce the transport burden before learning begins, while the perturbation scale $\sigma$ explicitly controls the injected noise energy. A large $\sigma$ increases the noise term and can dilute the benefit of the structured prior. In practice, an overly small $\sigma$ may also reduce stochastic regularization and make the model more dependent on deterministic prior bias, motivating the empirical tuning of $\sigma$.

## A.2. Effect of Perturbation Scale

The target-energy decomposition suggests that the perturbation scale $\sigma$ controls the trade-off between preserving the boundary-informed prior and injecting stochastic regularization. To empirically examine this effect, we sweep $\sigma \in \{0.01, 0.05, 0.10, 0.15, 0.20, 0.30\}$ and report gap-averaged position MAE across 6h, 8h, and 10h gaps over three random seeds.

Across both linear and Hermite priors and across Euler and Dopri5 solvers, performance is consistently best around $\sigma = 0.10$. When $\sigma$ is too small, the model may become overly tied to the deterministic prior; when $\sigma$ is too large, the structured source is increasingly dominated by injected noise. This supports the interpretation that moderate perturbation preserves useful boundary structure while providing sufficient stochastic regularization.

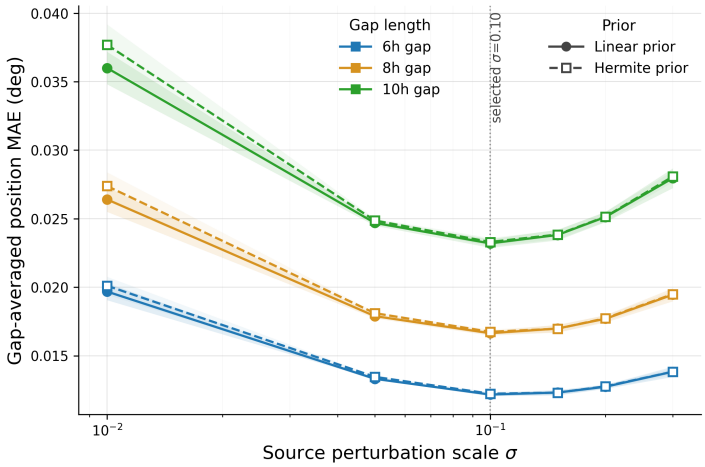


*Figure 6.* Effect of the source perturbation scale $\sigma$. Each value reports gap-averaged position MAE over 6h, 8h, and 10h gaps across three random seeds. Moderate perturbation around $\sigma = 0.10$ gives the best performance, while very small or large perturbations degrade accuracy.

## B. Implementation Details

### B.1. Computational Environment

All experiments are conducted on an Ubuntu 24.04.4 LTS server equipped with two AMD EPYC 7502 32-Core CPUs, 2.0 TiB RAM, and four NVIDIA GeForce RTX 3090 GPUs with 24 GB memory each.

### B.2. Architecture and Hyperparameters

*Table 5.* AISFlow architecture configuration.

| Component | Configuration |
|---|---|
| Input window | 144 timesteps, 4 channels |
| Observation encoder | 3-layer Transformer |
| Hidden size | 128 |
| Attention heads | 4 |
| Velocity network | 1D U-Net |
| Base channels | 128 |
| Channel multipliers | $(1, 2, 4)$ |
| Total parameters | 17.36M |

*Table 6.* Training and inference hyperparameters.

| Hyperparameter | Value |
|---|---|
| Optimizer | Adam |
| Learning rate | $2 \times 10^{-4}$ |
| Batch size | 256 |
| Epochs | 200 |
| Gradient clipping | 1.0 |
| Source perturbation | $\sigma = 0.10$ |
| Ensemble size | $N = 16$ |
| ODE solver | Dopri5 |