

Beyond Observations: Reconstruction Error-Guided Irregularly Sampled Time Series Representation Learning

Jiexi Liu^{1*}, Meng Cao^{2,3*}, Songcan Chen^{2,3†}

¹School of Computer and Artificial Intelligence, Nanjing University of Finance and Economics

²College of Computer Science and Technology, Nanjing University of Aeronautics and Astronautics

³MITT Key Laboratory of Pattern Analysis and Machine Intelligence

{liujiexi, meng.cao, s.chen}@nuaa.edu.cn

Abstract

Irregularly sampled time series (ISTS), characterized by non-uniform time intervals with natural missingness, are prevalent in real-world applications. Existing approaches for ISTS modeling primarily rely on observed values to impute unobserved ones or infer latent dynamics. However, these methods overlook a critical source of learning signal: *the reconstruction error inherently produced during model training*. Such error implicitly reflects how well a model captures the underlying data structure and can serve as an informative proxy for unobserved values. To exploit this insight, we propose **iTimER**, a simple yet effective self-supervised pre-training framework for ISTS representation learning. iTimER models the distribution of reconstruction errors over observed values and generates pseudo-observations for unobserved timestamps through a mixup strategy between sampled errors and the last available observations. This transforms unobserved timestamps into *noise-aware training targets*, enabling meaningful reconstruction signals. A Wasserstein metric aligns reconstruction error distributions between observed and pseudo-observed regions, while a contrastive learning objective enhances the discriminability of learned representations. Extensive experiments on classification, interpolation, and forecasting tasks demonstrate that iTimER consistently outperforms state-of-the-art methods under the ISTS setting.

Introduction

Irregularly sampled time series (ISTS) are ubiquitous in real-world applications, ranging from healthcare (Goldberger et al. 2000; Reyna et al. 2020; Liu, Cao, and Chen 2025), meteorology (Schulz and Stattegger 1997; Cao et al. 2018) to transportation (Chen et al. 2022; Tang et al. 2020). Arising from factors such as sensor failures, transmission distortions, or cost-driven acquisition policies, ISTS exhibit inconsistent intervals between consecutive sampling timestamps within a variable, asynchronous sampling across multiple variables and sometimes sampling sparsity. As shown in Fig.1, these characteristics manifest as **natural missingness**, i.e., unobserved values occurring, violating the assumption of a coherent, fixed-dimensional feature space that

*These authors contributed equally.

†Corresponding Author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

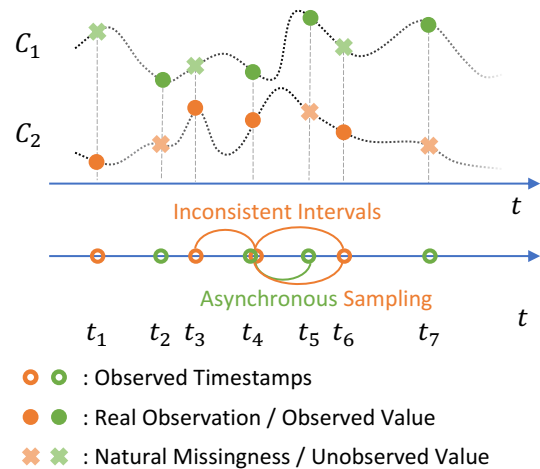


Figure 1: Multivariate irregularly sampled time series (ISTS) with two variables C_1 and C_2 , exhibiting non-uniform sampling intervals with natural missingness.

underlies most conventional time series analysis methods (Shukla and Marlin 2020).

Recent studies have explored multiple directions to address the above challenge and model temporal dependencies with missing patterns. A common line of work leverages explicit imputation of missing values as a preprocessing step before learning representations or performing downstream tasks (Che et al. 2018; Camino, Hammerschmidt, and State 2019; Tashiro et al. 2021; Zhang et al. 2021b; Chen et al. 2022; Fan 2022; Du, Côté, and Liu 2023; Fons et al. 2025). However, such direct imputation can be unreliable under high missingness ratio and may introduce noise or bias (Zhang et al. 2021a; Wu et al. 2021; Agarwal et al. 2023; Sun et al. 2024). Therefore, other kind of method proposes end-to-end architectures that typically model the data using continuous-time dynamics (Kidger et al. 2020; Rubanova, Chen, and Duvenaud 2019; Jhin et al. 2022; Oh et al. 2025; Zhang et al. 2025) or learning discrete regularly-sampled representations (Shukla and Marlin 2021; Zhang et al. 2023; Yalavarthi et al. 2024; Agarwal et al. 2023; Zhang et al. 2024b). In parallel, a growing body of work adopts self-supervised learning paradigms for ISTS, where

supervision is derived from the data itself (Chowdhury et al. 2023; Beebe-Wang et al. 2023). These methods apply contrastive learning to align representations of different augmented views, while others design reconstruction-based objectives from perturbed dropout inputs to encourage informative representations.

Despite the above advancements, existing approaches for learning from ISTS exhibit a fundamental limitation: they predominantly rely on the observed data as the only available source of training signals, except for label information. As a result, the model’s behavior in unobserved regions remains underconstrained, and its representational potential is largely overlooked. Crucially, these approaches tend to disregard model-intrinsic signals that arise naturally during optimization. Among them, the reconstruction error that quantifies the discrepancy between inputs and outputs encodes implicit information about the model’s current understanding of the data. While self-supervised masked modeling methods usually minimize such error over observed regions, they typically treat it as a loss term, rather than a source of feedback that can inform learning beyond the observations. In our opinion, the reconstruction error itself carries rich statistical cues that can be leveraged to guide learning in unobserved regions. In particular, the error distribution reflects the model’s uncertainty and inductive bias, offering a unique opportunity to synthesize auxiliary training signals without relying on external supervision information or heuristic augmentations. Motivated by this insight, in this work, we propose **iTimER**, a novel self-supervised pretraining framework that treats reconstruction error as an informative learning signal. By modeling its distribution and propagating it to unobserved timestamps, we construct pseudo-observations that enable distributionally aligned learning. This strategy not only enriches the training signal but also paves the way for more reliable and expressive representations in ISTS without requiring any labeled data.

Rather than merely imputing missingness or injecting arbitrary noise, iTimER leverages the model’s own reconstruction behavior as a proxy for learning uncertainty. Specifically, we estimate a distribution of reconstruction errors over the observed timestamps and propagate this signal to unobserved regions via a sampling-and-mixup mechanism. This strategy reflects both the model’s uncertainty and its inductive preferences, enabling the generation of pseudo-observations aligned with the model’s learned reconstruction error distribution. As observations in ISTS are not uniformly distributed over time but are often concentrated around specific events or high-frequency periods, leaving other segments completely unobserved. As a result, the observed data only provide a partial view of the series, leading models trained on these inputs to learn biased dynamics that can impair generalization and downstream performance. Although some recent methods (Tashiro et al. 2021; Islam, Tadepalli, and Fern 2025) attempt to mitigate this issue by adding and removing noise during training, they typically overlook the uneven, biased nature of real-world sampling. In contrast, iTimER constructs pseudo-observations based on the model’s own uncertainty, enabling a noise-aware learning signal. Rather than treating the observed data

as perfect learning signals, iTimER embraces its imperfections and uses them as a proxy to guide representation learning for ISTS.

To further ensure consistency between real and pseudo signals, iTimER imposes a Wasserstein metric constraint between the reconstruction error distributions of observed and pseudo-observed regions. This encourages the model to maintain similar uncertainty patterns across both parts of the sequence, effectively mitigating the impact of sampling bias. Additionally, iTimER employs contrastive learning to enhance the discriminative power of the learned representations, enabling the encoder to better capture meaningful variations in temporal dynamics.

To sum up, iTimER uses self-generated, noise-aware supervision signals that reveal how the model generalizes beyond the observed data to learn meaningful representations under ISTS conditions. Our main contributions are summarized as follows:

- We uncover the untapped value of reconstruction error as an informative self-supervised signal, providing a novel perspective for representation learning in ISTS.
- We enforce consistency between real and pseudo-observation series by aligning their reconstruction error distributions, improving the reliability for better representation learning.
- We present a task-agnostic pre-training model for ISTS, applicable to various downstream tasks, including classification, forecasting, and interpolation.

Related Work

Representation Learning for Irregularly Sampled Time Series

Recent advances in ISTS analysis have increasingly shifted toward learning expressive and discriminative representations. A common kind of imputation-based methods impute unobserved values from real observations before downstream tasks (Che et al. 2018; Chen et al. 2022; Fan 2022; Du, Côté, and Liu 2023). However, inappropriate imputation can introduce misleading noise or structural bias, especially under sparse observations that negatively impact model performance (Zhang et al. 2021a; Agarwal et al. 2023). While end-to-end models aggregate observation values for each variable to get discrete hidden states using attention mechanisms (Shukla and Marlin 2021; Zhang et al. 2023; Zhou et al. 2025), Graph Neural Networks (GNNs) (Zhang et al. 2021a; Yalavarthi et al. 2024; Zhang et al. 2024b), and Recurrent Neural Networks (RNNs) (De Brouwer et al. 2019; Schirmer et al. 2022; Agarwal et al. 2023) -based methods. Such methods struggle to capture continuous temporal dependencies. To address this, Neural ODE-based approaches (Kidger et al. 2020; Rubanova, Chen, and Duvenaud 2019; Jhin et al. 2022; Oh et al. 2025; Zhang et al. 2025) have been proposed to model ISTS in continuous time.

Recent work has explored self-supervised learning for ISTS representation learning by deriving supervision directly from the data itself (Chowdhury et al. 2023; Beebe-Wang et al. 2023). Self-supervised methods design pretext

tasks that encourage the model to learn meaningful representations without relying on downstream labels, including contrastive learning that constructs augmented views from real observations for contrastive learning and masked modeling methods that drop parts of the input data, and predict the dropout values by minimizing the reconstruction error between the predicted and original values. When facing ISTS, the non-uniform time intervals with natural missingness make it difficult to construct meaningful positive pairs and waste valuable real observations and potentially distort the learning signal.

Different from existing masked modeling methods, iTimER leverages reconstruction error not just as a loss, but as a learning signal to model uncertainty. It estimates error distributions from observed data and synthesizes pseudo-observations, enabling noise-aware representation learning without explicit imputation or artificial noise (Tashiro et al. 2021; Islam, Tadepalli, and Fern 2025).

Reconstruction Error as a Learning Signal

Across domains, reconstruction error, also referred to as residual, has proven to be more than a loss metric but serves as a valuable signal for model guidance. In image inpainting, Yu et al. (2018) employs reconstruction error in a coarse-to-fine architecture, where the residuals from the coarse stage are passed to the fine stage as structural cues, improving detail consistency in the final output. Similarly, in robust subspace learning, Meng and De La Torre (2013) embeds reconstruction error into the optimization loop to jointly learn the data structure and noise characteristics. In video anomaly detection, works such as Hasan et al. (2016) and Liu et al. (2018) exploit spatially localized residuals to identify abnormal events.

These examples suggest that reconstruction error can highlight uncertainty, noise, or outliers and serve as an informative learning cue. iTimER builds on this idea by modeling the distribution of reconstruction error over observed regions in ISTS, then sampling from this distribution to generate uncertainty-aware pseudo-observations. These are used not as direct labels, but as a proxy signal that drives both contrastive representation learning and distributional consistency. To our knowledge, iTimER is the first to formalize this strategy in the ISTS domain, using reconstruction error not as a penalty, but as a means of self-supervised guidance.

Proposed iTimER Framework

Problem Setup and Method Outline

Let $\mathcal{X} = \{\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_N\}$ denote a dataset of N ISTS instances, where each instance $\mathbf{X}_n \in \mathbb{R}^{T \times C}$ contains C variables observed over T time steps. The binary mask $\mathbf{M}_n \in \{0, 1\}^{T \times C}$ is a missingness indicator for \mathbf{X}_n , where $m_{n,t,c} = 0$ indicates there is no observation at timestamp t in variable c and $m_{n,t,c} = 1$ means it is observed. Our objective is to learn a nonlinear encoding function $f_\theta : \mathbb{R}^{T \times C} \rightarrow \mathbb{R}^{\tau \times D}$ that maps each ISTS instance into a representation $\mathbf{R}_n = f_\theta(\mathbf{X}_n)$, where τ is the length of the representation and D is the dimension. *Here, for brevity, we omit the data*

case index n for the n -th instance and c for the c -th variable in an instance when the context is clear.

Specifically, as illustrated in Fig.2, given an ISTS instance \mathbf{X} with its missingness indicator \mathbf{M} , we employ an encoder-decoder architecture, where f_θ maps the input series to a latent representation \mathbf{R} and g_ϕ reconstructs the input from it to obtain reconstruction $\hat{\mathbf{X}}$. We first estimate the reconstruction error distribution $\epsilon_t \sim \mathcal{N}(\mu_\epsilon, \sigma_\epsilon^2)$ from $x_t - \hat{x}_t$, where $m_t = 1$. Then, for unobserved timestamps, i.e., $m_t = 0$, we sample errors $\tilde{\epsilon}_t$ from $\mathcal{N}(\mu_\epsilon, \sigma_\epsilon^2)$ and mixup it with the last available observation x_{t-1} to obtain the pseudo-observation \tilde{x}_t . This results finally becoming the pseudo-observation series $\tilde{\mathbf{X}}$ that preserves the original timestamps but contains pseudo values in place of the unobserved timestamps.

By jointly optimizing distribution alignment, representation consistency and reconstruction, this training strategy guides the encoder toward learning more informative and discriminative representations \mathbf{R} for downstream tasks, and we will elaborate on the key steps in the following sections.

Pseudo-Observation Learning

Modeling Reconstruction Uncertainty. For any instance \mathbf{X} , to model the uncertainty of reconstruction, we begin by reconstructing the observed timestamps, i.e., mask $m_t = 1$ by using an encoder-decoder model $\hat{x}_t = g_\phi(f_\theta(x_t))$. The reconstruction error is then computed as $\epsilon_t = x_t - \hat{x}_t$. To capture its statistical properties, we assume the error distribution ϵ_t following a Gaussian distribution $\epsilon_t \sim \mathcal{N}(\mu_\epsilon, \sigma_\epsilon)$, which can provide a basis for generating plausible pseudo-observations for the unobserved timestamps.

We leverage a Gaussian form for the reconstruction error distribution, which aligns with the empirical observation that residuals from deep time series models tend to follow near-normal patterns. This is also suitable for ISTS (Shukla and Marlin 2021; Fortuin et al. 2020), where the stochastic nature of sampling times introduces non-uniform uncertainty. A Gaussian approximation allows us to quantify such uncertainty and synthesize pseudo-observations in the absence of complete data.

In addition, to preserve informative signals while maintaining stability, a momentum strategy is adopted to utilize historical information through a momentum coefficient ρ for the reconstruction error as

$$\mu_\epsilon^h = \rho \cdot \mu_\epsilon^{h-1} + (1-\rho) \cdot \mu_\epsilon, \quad \sigma_\epsilon^h = \rho \cdot \sigma_\epsilon^{h-1} + (1-\rho) \cdot \sigma_\epsilon, \quad (1)$$

where μ_ϵ^{h-1} and σ_ϵ^{h-1} denote the mean and standard deviation at the $(h-1)$ -th iteration, separately.

Error Sampling at Unobserved Timestamps. For unobserved timestamps, i.e., $m_t = 0$, we synthesize pseudo-observations by mixing a sampled reconstruction error with a temporal contextual anchor value, such as the global mean, local moving average, or the last observation value. In this paper, we choose the last observation as it can better preserve temporal continuity and aligns with the causal nature of time series. Compared to the global mean or local moving average, it is not only a simple processing, but also faithfully

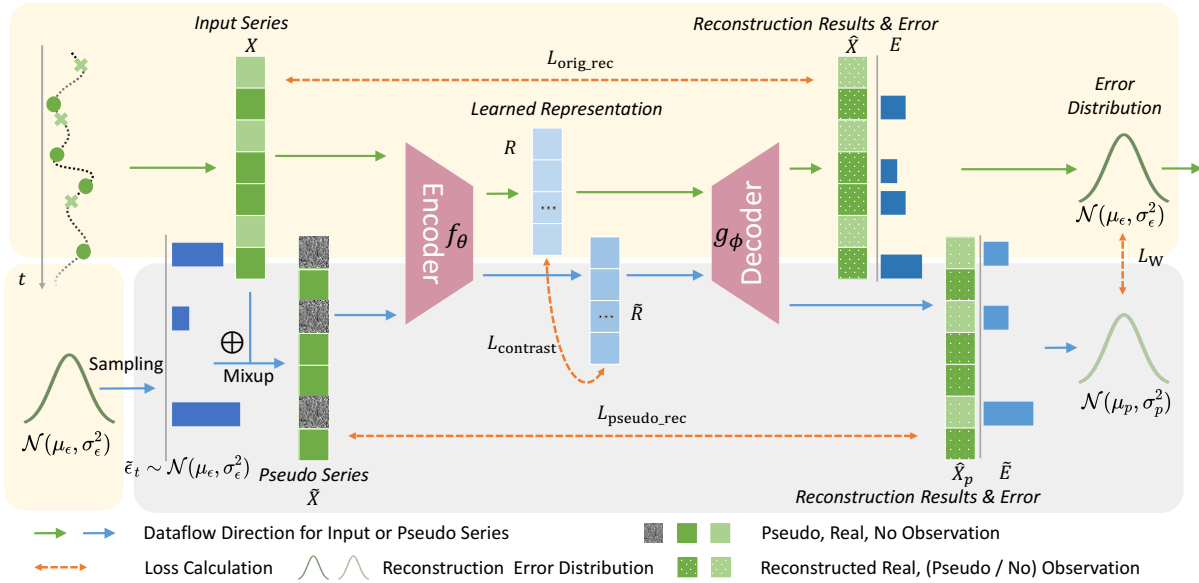


Figure 2: **The iTIMER framework.** iTIMER leverages reconstruction error from the original sequence to generate pseudo observations via mixup uncertainty-aware sampling and the last observed value. Both real and pseudo-observation series are encoded and reconstructed, with consistency enforced at both representation and error distribution levels.

maintains local dynamics. Then, we obtain the Mixup strategy defined as

$$\tilde{x}_t = \alpha_t \cdot \bar{x} + (1 - \alpha_t) \cdot \tilde{\epsilon}_t, \quad \tilde{\epsilon}_t \sim \mathcal{N}(\mu_\epsilon^h, (\sigma_\epsilon^h)^2) \quad (2)$$

where $\alpha_t \in [0, 1]$ controls the mixing ratio. This approach generates softly regularized imputations that balance between global structure and instance-specific uncertainty. In ISTS, where temporal gaps vary significantly and structural assumptions may not hold, such a mixup-based generation helps avoid overconfident or abrupt imputations, encourages smoother continuity, and enhances representation learning by exposing the model to more realistic distributional variations.

Constructing Complete Pseudo-Observation Series. To form a complete pseudo-observation series, we integrate the sampled values with the original observations based on the observed mask. Specifically, for each time step t , if the original value x_t is observed (i.e., $m_t = 1$), we retain the original observation. If x_t is missing (i.e., $m_t = 0$), we replace it with the generated pseudo-observation \tilde{x}_t introduced above. Formally, the constructed series is defined as:

$$\tilde{x}_t = \begin{cases} x_t, & \text{if } m_t = 1 \\ \tilde{x}_t, & \text{if } m_t = 0 \end{cases} \quad (3)$$

This procedure ensures that the final sequence preserves all original valid observations while providing realistic, distribution-aware proxy for the missingness, enabling the model to learn from a more complete and semantically coherent input. \tilde{X} then follows the similar procedure that reconstructed by $g_\phi(f_\theta(\tilde{X}))$ and obtain the reconstruction error distribution $\mathcal{N}(\mu_p, \sigma_p^2)$ where $m_t = 0$.

Learning Objective Design

Our main step of iTIMER is to enforce reconstruction error consistency. The unobserved timestamps in ISTS have no ground truth and exhibit non-uniform sampling intervals that cannot directly supervise the model learning in those regions. By matching the statistical profile of errors between observed and imputed points, we provide a principled learning signal on the unobserved timestamps that 1) quantifies the model’s uncertainty in unobserved timestamps, 2) prevents overconfident, biased imputations, and 3) encourages the encoder to learn representations that generalize across both observed and missing parts. This error distribution alignment thus serves as a proxy for true reconstruction quality where observations are unavailable, leading to uncertainty-aware ISTS representations.

Here, to match the reconstruction error distributions of the real and pseudo-observation series, we use the 2-Wasserstein distance (Panaretos and Zemel 2019) between $P_r = \mathcal{N}(\mu_\epsilon, \sigma_\epsilon^2)$ and $P_p = \mathcal{N}(\mu_p, \sigma_p^2)$ as follows

$$\begin{aligned} L_W &= D_{2\text{-Wasserstein}}(P_r \parallel P_p) \\ &= \|\mu_\epsilon - \mu_p\|^2 + \|\sigma_\epsilon - \sigma_p\|^2 \end{aligned} \quad (4)$$

Moreover, to ensure robust and informative representation learning under ISTS, we incorporate both contrastive learning and dual reconstruction objectives. The contrastive loss encourages the latent representations of the original and pseudo-observation series to remain close, thereby mitigating representation drift caused by missing values. This alignment in the embedding space not only enhances the semantic consistency between observed and reconstructed data but also facilitates the learning of stable, discriminative features that are less sensitive to sampling irregularities or noise.

Input ISTS \mathbf{X} and pseudo-observation series $\widetilde{\mathbf{X}}$ are encoded using a shared encoder $f_\theta(\cdot)$, producing latent representations \mathbf{R} and $\widetilde{\mathbf{R}}$, respectively:

$$\mathbf{R} = f_\theta(\mathbf{X}), \quad \widetilde{\mathbf{R}} = f_\theta(\widetilde{\mathbf{X}}) \quad (5)$$

Therefore, the contrastive loss function is:

$$L_{\text{contrast}} = - \sum_{i=1}^{|\mathcal{B}|} \log \frac{\exp(\mathbf{R}_i \cdot \widetilde{\mathbf{R}}_i)}{\sum_{j=1}^{|\mathcal{B}|} (\mathbf{R}_i \cdot \widetilde{\mathbf{R}}_i) + \mathbb{I}_{[i \neq j]}(\mathbf{R}_i \cdot \mathbf{R}_j)} \quad (6)$$

where the \mathbb{I} is the indicator function, $|\mathcal{B}|$ indicates batch size.

The dual reconstruction losses Eq.(7), one on the original observations and the other on the pseudo-observation series, serve complementary roles. The original reconstruction loss serves as a fundamental self-supervision signal to ensure that the encoder f_θ and decoder g_ϕ can effectively capture the underlying structure of the observed data and accurately reconstruct the ISTS. While the pseudo-observation reconstruction loss evaluates whether the model can still produce stable and reasonable reconstructions. It strengthens the model’s ability to handle pseudo-observations, helping it learn the underlying distribution of the imputed regions.

$$\begin{aligned} L_{\text{orig_rec}} &= \|\mathbf{M} \odot (\mathbf{X} - \widehat{\mathbf{X}})\|^2 \\ L_{\text{pseudo_rec}} &= \|\mathbf{M} \odot (\widetilde{\mathbf{X}} - \widehat{\mathbf{X}}_p)\|^2 \end{aligned} \quad (7)$$

The overall loss function is a weighted combination of the above terms:

$$L = \alpha L_W + \beta L_{\text{contrast}} + \frac{1}{2}(L_{\text{orig_rec}} + L_{\text{pseudo_rec}}) \quad (8)$$

where α and β are hyperparameters controlling the relative importance of each term.

Experiments

In this section, we show the effectiveness of the iTimER framework across 3 mainstream time series downstream tasks, including classification, interpolation, and forecasting. The results are reported as mean and standard deviation values calculated over 5 independent runs. We use Time-Feature Attention (TFA) and Feature-Feature Attention (FFA) as the backbone encoder (Chowdhury et al. 2023). The **bold** font highlights the top-performing method, while the underlined text marks the runner-up. Additional experimental setup and encoder details are provided in the Appendix due to space constraints.

Experimental Setup

Datasets and Metrics. For the *Classification* task, we evaluate on three real-world ISTS datasets from healthcare and human activity domains: (1) **P19** (Reyna et al. 2020) includes 38,803 ICU patients monitored by 34 sensors, with a missing ratio of 94.9%. (2) **P12** (Goldberger et al. 2000) contains temporal measurements from 11,988 patients over the first 48 hours of ICU stay, using 36 sensors and a missing ratio of 88.4%. (3) **PAM** (Reiss and Stricker 2012) consists of 5,333 sequences of 8 daily activities measured by 17 sensors, with 60.0% missingness.

P19 and P12 are imbalanced binary label datasets while PAM dataset is balance containing 8 classes. Following the standard practice and prior works (Li, Li, and Yan 2023; Zhang et al. 2021a), we split each dataset into training/validation/test sets using an 8 : 1 : 1 ratio and fixed indices for fair comparison. For the *imbalanced* P12 and P19 datasets, we report Area Under the Receiver Operating Characteristic Curve (AUROC) and Area Under the Precision-Recall Curve (AUPRC), while for the nearly balanced PAM dataset, we use Accuracy, Precision, Recall, and F1 Score. Higher values across all these metrics indicate better performance.

For *Interpolation and Forecasting* tasks, we use three natural ISTS datasets from the healthcare and activity domains: (1) **Physionet** (Goldberger et al. 2000) includes 12,000 ICU patients monitored by 41 sensors, with a missing ratio of 85.7%. (2) **MIMIC** (Johnson et al. 2016) contains 23,457 ISTS samples covering the first 48 hours of ICU stays, with 96 variables and 96.7% missingness. (3) **Human Activity** (Frank, Asuncion et al. 2010) comprises 5,400 samples of 3D positional data collected from 12 sensors, with 75.0% missingness.

We randomly split each dataset into training, validation, and test sets with a 6 : 2 : 2 ratio, applying min-max normalization to the original observation values. For both tasks, we evaluate performance using Mean Squared Error (MSE) and Mean Absolute Error (MAE), where lower values indicate better performance. Complete results are reported in the appendix.

Baselines. To evaluate the performance of iTimER in ISTS *Classification*, we incorporate the following baseline models for a fair comparison, including Transformer (Vaswani et al. 2017), MTGNN (Wu et al. 2020), DGM²-O (Wu et al. 2021), IP-Net (Shukla and Marlin 2018), GRU-D (Che et al. 2018), SeFT (Horn et al. 2020), mTAND (Shukla and Marlin 2021), Raindrop (Zhang et al. 2021a), Warpformer (Zhang et al. 2023) and ViTST (Li, Li, and Yan 2023). Specially, we also compare with several Pre-trained Language Model (PLM)-based methods for regularly sampled time series analysis, such as FPT (Zhou et al. 2023) and Time-LLM (Jin et al. 2024), as well as pre-training ISTS models PrimeNet (Chowdhury et al. 2023).

For ISTS *Interpolation and Forecasting* tasks, except adapting the representative baselines above to these two tasks, we further incorporate several models designed for the ISTS prediction tasks, including Latent-ODE (Rubanova, Chen, and Duvenaud 2019), Neural Flow (Biloš et al. 2021), CRU (Schirmer et al. 2022), t-PatchGNN (Zhang et al. 2024b) and ISTS-PLM (Zhang et al. 2024a), a PLM-based method for ISTS.

Main Results

As shown in Table 1, iTimER consistently achieves strong performance, demonstrating its effectiveness for ISTS classification tasks. Specifically, in the *imbalanced* binary classification tasks, iTimER outperforms prior methods on the P12 dataset by clear margins in both AUROC and AUPRC, indicating greater sensitivity to rare yet informative signals in sparse ICU data. It maintains competitive results on

Method	P12		P19		PAM			
	AUROC	AUPRC	AUROC	AUPRC	Accuracy	Precision	Recall	F1 score
Transformer	83.3 ± 0.7	47.9 ± 3.6	80.7 ± 3.8	42.7 ± 7.7	83.5 ± 1.5	84.8 ± 1.5	86.0 ± 1.2	85.0 ± 1.3
MTGNN	74.4 ± 6.7	35.5 ± 6.0	81.9 ± 6.2	39.9 ± 8.9	83.4 ± 1.9	85.2 ± 1.7	86.1 ± 1.9	85.9 ± 2.4
DGM ² -O	84.4 ± 1.6	47.3 ± 3.6	86.7 ± 3.4	44.7 ± 11.7	82.4 ± 2.3	85.2 ± 1.2	83.9 ± 2.3	84.3 ± 1.8
IP-Net	82.6 ± 1.4	47.6 ± 3.1	84.6 ± 1.3	38.1 ± 3.7	74.3 ± 3.8	75.6 ± 2.1	77.9 ± 2.2	76.6 ± 2.8
GRU-D	81.9 ± 2.1	46.1 ± 4.7	83.9 ± 1.7	46.9 ± 2.1	83.3 ± 1.6	84.6 ± 1.2	85.2 ± 1.6	84.8 ± 1.2
SeFT	73.9 ± 2.5	31.1 ± 4.1	81.2 ± 2.3	41.9 ± 3.1	67.1 ± 2.2	70.0 ± 2.4	68.2 ± 1.5	68.5 ± 1.8
mTAND	84.2 ± 0.8	48.2 ± 3.4	84.4 ± 1.3	50.6 ± 2.0	92.9 ± 0.8	93.8 ± 0.8	94.0 ± 0.9	93.8 ± 0.8
Raindrop	82.8 ± 1.7	44.0 ± 3.0	87.0 ± 2.3	51.8 ± 5.5	88.5 ± 1.5	89.9 ± 1.5	89.9 ± 0.6	89.8 ± 1.0
Warpformer	83.4 ± 0.9	47.2 ± 3.7	<u>88.8</u> ± 1.7	55.2 ± 3.9	94.3 ± 0.6	95.8 ± 0.8	94.8 ± 1.0	95.2 ± 0.6
ViTST	<u>85.1</u> ± 0.8	<u>51.1</u> ± 4.1	89.2 ± 2.0	<u>53.1</u> ± 3.4	<u>95.8</u> ± 1.3	<u>96.2</u> ± 1.3	<u>96.1</u> ± 1.1	<u>96.5</u> ± 1.2
FPT	84.8 ± 1.1	50.7 ± 3.0	87.3 ± 2.9	51.6 ± 3.6	94.0 ± 1.4	95.3 ± 0.9	94.7 ± 1.1	94.9 ± 1.1
Time-LLM	84.4 ± 1.8	50.2 ± 1.6	85.1 ± 2.6	50.1 ± 3.4	93.4 ± 1.2	94.2 ± 1.3	94.7 ± 1.0	94.4 ± 1.1
PrimeNet	84.9 ± 0.6	49.8 ± 2.7	84.4 ± 1.3	39.7 ± 3.1	95.3 ± 0.5	96.1 ± 0.3	95.5 ± 0.6	95.7 ± 0.4
iTimER	85.7 ± 0.8	52.0 ± 2.1	87.1 ± 0.6	45.6 ± 3.5	96.1 ± 0.8	96.7 ± 0.5	96.4 ± 0.9	96.6 ± 0.7

Table 1: Overall performance comparison on ISTS *Classification* task.

Method	PhysioNet	MIMIC	Human Activity
	MSE × 10 ⁻³	MSE × 10 ⁻²	MSE × 10 ⁻³
GRU-D	6.18 ± 0.23	2.06 ± 0.05	2.74 ± 0.09
SeFT	9.46 ± 0.12	2.12 ± 0.02	14.95 ± 0.03
Raindrop	10.65 ± 0.12	2.31 ± 0.04	15.21 ± 0.12
Warpformer	6.37 ± 0.34	1.93 ± 0.06	2.59 ± 0.15
mTAND	5.65 ± 0.08	1.93 ± 0.05	2.07 ± 0.17
Latent-ODE	6.84 ± 0.34	1.89 ± 0.08	3.12 ± 0.22
Neural Flow	6.77 ± 0.06	2.18 ± 0.11	3.73 ± 0.06
CRU	10.30 ± 0.10	2.52 ± 0.04	7.17 ± 0.32
t-PatchGNN	4.75 ± 0.03	1.55 ± 0.08	1.95 ± 0.12
FPT	12.24 ± 0.07	3.71 ± 0.01	2.85 ± 0.09
Time-LLM	12.43 ± 0.08	3.63 ± 0.05	2.92 ± 0.01
ISTS-PLM	<u>4.55</u> ± 0.08	<u>1.47</u> ± 0.01	<u>1.93</u> ± 0.01
iTimER	2.86 ± 0.04	0.13 ± 0.00	1.82 ± 0.01

Table 2: Overall performance comparison on ISTS *Interpolation* task.

P19 and achieves SOTA performance on the balanced PAM dataset, demonstrating its consistent effectiveness across varying levels of class imbalance and missingness. Moreover, as shown in Fig. 3, iTimER achieves this without sacrificing efficiency. Its favorable balance between performance and complexity makes it especially suitable for ISTS.

Table 2 presents the performance comparison for ISTS interpolation tasks, where 30% of the observation timestamps are randomly masked and the model is tasked with reconstructing them using unmasked data. iTimER consistently achieves the best performance across all three datasets, outperforming recent SOTA methods. Its superior results, particularly on high-missingness datasets like MIMIC and PhysioNet, indicate a strong ability to capture local structure and uncertainty without relying on explicit imputation or continuous-time solvers.

Table 3 shows the performance on ISTS forecasting tasks,

Method	PhysioNet	MIMIC	Human Activity
	MSE × 10 ⁻³	MSE × 10 ⁻²	MSE × 10 ⁻³
GRU-D	5.59 ± 0.09	1.76 ± 0.03	2.94 ± 0.05
SeFT	9.22 ± 0.18	1.87 ± 0.01	12.20 ± 0.17
Raindrop	9.82 ± 0.08	1.99 ± 0.03	14.92 ± 0.14
Warpformer	5.94 ± 0.35	1.73 ± 0.04	2.79 ± 0.04
mTAND	6.23 ± 0.24	1.85 ± 0.06	3.22 ± 0.07
Latent-ODE	6.05 ± 0.57	1.89 ± 0.19	3.34 ± 0.11
Neural Flow	7.20 ± 0.07	1.87 ± 0.05	4.05 ± 0.13
CRU	8.56 ± 0.26	1.97 ± 0.02	6.97 ± 0.78
t-PatchGNN	4.98 ± 0.08	1.69 ± 0.03	2.66 ± 0.03
FPT	10.95 ± 0.02	4.00 ± 0.03	3.03 ± 0.09
Time-LLM	11.56 ± 0.19	4.41 ± 0.01	3.21 ± 0.01
ISTS-PLM	4.92 ± 0.05	1.64 ± 0.02	2.58 ± 0.03
iTimER	3.64 ± 0.05	0.14 ± 0.00	2.75 ± 0.03

Table 3: Overall performance comparison on ISTS *Forecasting* task.

where we follow the common setup in Zhang et al. (2024b) that uses the first 24 hours of data to predict the next 24 hours on PhysioNet and MIMIC datasets, and uses the first 3,000 ms to predict the next 1,000 ms on Human Activity data. iTimER achieves SOTA results on PhysioNet and MIMIC, and ranks among the top methods on Human Activity. These results suggest that iTimER’s reconstruction-aware representation learning generalizes well to long-range extrapolation, even under significant sparsity and distribution shift.

In nearly all cases, iTimER exhibits consistently *low standard deviation*, indicating it is a reliable and robust model. Its performance remains stable across varying data samples and initial conditions, highlighting its strong potential to generalize well to new, unseen data. This stability and predictability are particularly valuable in sensitive domains such as medical diagnosis, where accurate and reliable pre-

Method	P12	PhysioNet-I	PhysioNet-F
	AUROC \uparrow	MSE $\times 10^{-3}$ \downarrow	MSE $\times 10^{-3}$ \downarrow
Baseline	84.6 \pm 2.1	5.03 \pm 0.04	6.10 \pm 0.29
Random	84.8 \pm 0.6	4.01 \pm 0.06	4.61 \pm 0.10
Constant	84.7 \pm 0.9	3.89 \pm 0.10	4.26 \pm 0.07
Only Error	85.1 \pm 0.8	3.08 \pm 0.06	3.77 \pm 0.05
Zero	85.0 \pm 1.9	3.06 \pm 0.10	3.94 \pm 0.12
Mean	85.4 \pm 1.0	2.91 \pm 0.06	3.66 \pm 0.08
MAve	85.8 \pm 1.0	2.83 \pm 0.01	3.64 \pm 0.02
w/o L_W	85.0 \pm 1.0	3.17 \pm 0.05	3.91 \pm 0.01
w/o L_{contrast}	85.4 \pm 0.7	3.02 \pm 0.05	3.78 \pm 0.03
iTimER	85.7 \pm 0.8	2.86 \pm 0.04	3.64 \pm 0.05

Table 4: Ablation results of iTimER on three downstream tasks, evaluating different pseudo-observation strategies and loss components (P12 for classification, PhysioNet-I for interpolation, and PhysioNet-F for forecasting).

dictions are essential in clinical settings.

Ablation Analysis and Efficiency Evaluation

Ablation Analysis. Table 4 presents the ablation study evaluating various pseudo-observation strategies and loss components in iTimER. Here, *Baseline* performs reconstruction-based pretraining only on the observed data, without generating pseudo-observations or modeling reconstruction errors for unobserved timestamps. It serves as a lower result that lacks uncertainty modeling and signal augmentation.

We first test substitutes for \tilde{x}_t , i.e., the proxy for the unobserved regions: Random samples (*Random*), a constant global mean (*Constant*), and using only the reconstruction error (*Only Error*). *Random* and *Constant* variants utilize uninformed or static values, leading to substantial degradation in all tasks, highlighting the importance of noise-aware generation. While *Only Error* variant relies completely on the reconstruction error component and performs better than the above two variants, but still underperforms compared to our full iTimER. This highlights that sampling from the error distribution $\mathcal{N}(\mu_\epsilon, \sigma_\epsilon^2)$ alone introduces meaningful inductive cues, along with the fact that without anchoring to observed values, it lacks the structural grounding necessary for robust representation learning.

Subsequently, in the mixup formulation $\tilde{x}_t = \alpha_t \cdot \bar{x} + (1 - \alpha_t) \cdot \tilde{\epsilon}_t$, we test three alternatives for \bar{x} : setting it to zero (*Zero*), the global mean (*Mean*), or a local moving average (*MAve*) with a window size of 5. Among them, the *MAve* strategy achieves the best results, suggesting that preserving local dynamics benefits pseudo-observation quality. However, it introduces additional computational cost due to windowed processing and potential extra parameters. Finally, removing either the Wasserstein loss or the contrastive consistency loss leads to clear performance drops, confirming that both are essential for aligning uncertainty and ensuring representational coherence. Overall, these results vali-

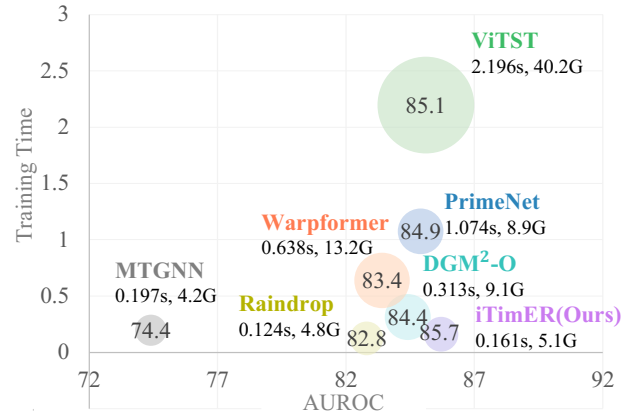


Figure 3: Efficiency comparisons in terms of Training Time (s) and Memory Usage (G) with the latest advanced models on the P12 datasets.

date our design choices and demonstrate that iTimER’s gains are from principled generation and multi-objective training.

Efficiency Evaluation. To demonstrate computational complexity, taking P12 in the classification task with a batch size of 50 as an example, iTimER achieves a training time cost of 0.161s per batch with 5.16 GB of memory usage. As in Fig.3, the closer a circle is to the **bottom-right corner** and the **smaller** its area, the higher the model’s classification accuracy, with faster training speed and lower memory usage. Our iTimER achieves lower time complexity than most other methods and uses much less memory, especially less than ViTST, which also performs well on classification tasks.

Although we achieve SOTA performance for the P12 dataset, our training time cost and memory usage are not optimal. This is primarily due to the additional overhead introduced by computing reconstruction error distributions and generating uncertainty-aware pseudo-observations during training. These steps enable iTimER to model ISTS more effectively, leading to improved performance across tasks despite the slight increase in resource consumption.

Conclusion

In this work, we propose iTimER, a novel framework for ISTS representation learning that leverages reconstruction error distributions as informative training signals. By generating pseudo-observations guided by learned error distributions and enforcing consistency across original and pseudo-observation series, iTimER effectively enriches the training signal without relying on labels or explicit imputation. Extensive experiments across three downstream tasks demonstrate its strong performance and broad applicability.

Looking forward, the learned reconstruction error distribution offers a promising plug-and-play module that can be integrated into a wide range of ISTS models. Future work will explore its utility for enhancing uncertainty modeling, guiding attention mechanisms, or improving robustness under distribution shift.

Acknowledgements

The main framework of this work was developed under the supervision of Professor Songcan Chen at Nanjing University of Aeronautics and Astronautics.

The authors wish to thank all the donors of the original datasets, and everyone who provided feedback on this work. This work is supported by the Key Program of NSFC under Grant No.62376126.

References

- Agarwal, R.; Sinha, A.; Prasad, D. K.; Clausel, M.; Horsch, A.; Constant, M.; and Coubez, X. 2023. Modelling Irregularly Sampled Time Series Without Imputation. *arXiv preprint arXiv:2309.08698*.
- Beebe-Wang, N.; Ebrahimi, S.; Yoon, J.; Arik, S. O.; and Pfister, T. 2023. PAITS: Pretraining and Augmentation for Irregularly-Sampled Time Series. *arXiv preprint arXiv:2308.13703*.
- Biloš, M.; Sommer, J.; Rangapuram, S. S.; Januschowski, T.; and Günnemann, S. 2021. Neural flows: Efficient alternative to neural ODEs. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, volume 34, 21325–21337.
- Camino, R. D.; Hammerschmidt, C. A.; and State, R. 2019. Improving missing data imputation with deep generative models. *arXiv preprint arXiv:1902.10666*.
- Cao, W.; Wang, D.; Li, J.; Zhou, H.; Li, L.; and Li, Y. 2018. Brits: Bidirectional recurrent imputation for time series. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, volume 31.
- Che, Z.; Purushotham, S.; Cho, K.; Sontag, D.; and Liu, Y. 2018. Recurrent neural networks for multivariate time series with missing values. *Scientific Reports*, 8(1): 1–12.
- Chen, X.; Zhang, C.; Zhao, X.-L.; Saunier, N.; and Sun, L. 2022. Nonstationary temporal matrix factorization for multivariate time series forecasting. *arXiv preprint arXiv:2203.10651*.
- Chowdhury, R. R.; Li, J.; Zhang, X.; Hong, D.; Gupta, R. K.; and Shang, J. 2023. Primenet: Pre-training for irregular multivariate time series. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 37.
- De Brouwer, E.; Simm, J.; Arany, A.; and Moreau, Y. 2019. GRU-ODE-Bayes: Continuous modeling of sporadically-observed time series. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, volume 32.
- Du, W.; Côté, D.; and Liu, Y. 2023. Saits: Self-attention-based imputation for time series. *Expert Systems with Applications*, 219: 119619.
- Fan, J. 2022. Dynamic Nonlinear Matrix Completion for Time-Varying Data Imputation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Fons, E.; Sztrajman, A.; El-Laham, Y.; Ferrer, L.; Vyetrenko, S.; and Veloso, M. 2025. LSCD: Lomb–Scargle Conditioned Diffusion for Time series Imputation. In *Proceedings of International Conference on Machine Learning (ICML)*.
- Fortuin, V.; Baranchuk, D.; Rätsch, G.; and Mandt, S. 2020. Gp-vae: Deep probabilistic time series imputation. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 1651–1661. PMLR.
- Frank, A.; Asuncion, A.; et al. 2010. UCI machine learning repository. URL <http://archive.ics.uci.edu/ml>, 15: 22.
- Goldberger, A. L.; Amaral, L. A.; Glass, L.; Hausdorff, J. M.; Ivanov, P. C.; Mark, R. G.; Mietus, J. E.; Moody, G. B.; Peng, C.-K.; and Stanley, H. E. 2000. PhysioBank, PhysioToolkit, and PhysioNet: components of a new research resource for complex physiologic signals. *Circulation*, 101(23): e215–e220.
- Hasan, M.; Choi, J.; Neumann, J.; Roy-Chowdhury, A. K.; and Davis, L. S. 2016. Learning temporal regularity in video sequences. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 733–742.
- Horn, M.; Moor, M.; Bock, C.; Rieck, B.; and Borgwardt, K. 2020. Set functions for time series. In *Proceedings of International Conference on Machine Learning (ICML)*.
- Islam, M. R. U.; Tadepalli, P.; and Fern, A. 2025. Self-attention-based Diffusion Model for Time-series Imputation in Partial Blackout Scenarios. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 39, 17564–17572.
- Jhin, S. Y.; Lee, J.; Jo, M.; Kook, S.; Jeon, J.; Hyeong, J.; Kim, J.; and Park, N. 2022. Exit: Extrapolation and interpolation-based neural controlled differential equations for time-series classification and forecasting. In *Proceedings of ACM Web Conference*, 3102–3112.
- Jin, M.; Wang, S.; Ma, L.; Chu, Z.; Zhang, J. Y.; Shi, X.; Chen, P.-Y.; Liang, Y.; Li, Y.-F.; Pan, S.; and Wen, Q. 2024. Time-LLM: Time Series Forecasting by Reprogramming Large Language Models. In *Proceedings of International Conference on Learning Representations (ICLR)*.
- Johnson, A.; Pollard, T. J.; Shen, L.; Lehman, L.-w. H.; Feng, M.; Ghassemi, M.; Moody, B.; Szolovits, P.; Celi, L. A.; and Mark, R. G. 2016. MIMIC-III, a freely accessible critical care database *Sci. Data*, 3(1): 1.
- Kidger, P.; Morrill, J.; Foster, J.; and Lyons, T. 2020. Neural controlled differential equations for irregular time series. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, 6696–6707.
- Li, Z.; Li, S.; and Yan, X. 2023. Time Series as Images: Vision Transformer for Irregularly Sampled Time Series. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*.
- Liu, J.; Cao, M.; and Chen, S. 2025. TimeCHEAT: A Channel Harmony Strategy for Irregularly Sampled Multivariate Time Series Analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 39.
- Liu, W.; Luo, W.; Lian, D.; and Gao, S. 2018. Future frame prediction for anomaly detection—a new baseline. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 6536–6545.
- Meng, D.; and De La Torre, F. 2013. Robust matrix factorization with unknown noise. In *Proceedings of the*

- IEEE conference on computer vision and pattern recognition (CVPR)*, 1337–1344.
- Oh, Y.; Kam, S.; Lee, J.; Lim, D.-Y.; Kim, S.; and Bui, A. 2025. Comprehensive review of neural differential equations for time series analysis. *arXiv preprint arXiv:2502.09885*.
- Panaretos, V. M.; and Zemel, Y. 2019. Statistical aspects of Wasserstein distances. *Annual review of statistics and its application*, 6(1): 405–431.
- Reiss, A.; and Stricker, D. 2012. Introducing a new benchmarked dataset for activity monitoring. In *Proceedings of International Symposium on wearable computers*, 108–109. IEEE.
- Reyna, M. A.; Josef, C. S.; Jeter, R.; Shashikumar, S. P.; Westover, M. B.; Nemati, S.; Clifford, G. D.; and Sharma, A. 2020. Early prediction of sepsis from clinical data: the PhysioNet/Computing in Cardiology Challenge 2019. *Critical Care Medicine*, 48(2): 210–217.
- Rubanova, Y.; Chen, R. T.; and Duvenaud, D. K. 2019. Latent ordinary differential equations for irregularly-sampled time series. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, volume 32.
- Schirmer, M.; Eltayeb, M.; Lessmann, S.; and Rudolph, M. 2022. Modeling irregular time series with continuous recurrent units. In *Proceedings of International Conference on Machine Learning (Proceedings of International Conference on Machine Learning (ICML))*, 19388–19405. PMLR.
- Schulz, M.; and Stattegger, K. 1997. SPECTRUM: Spectral analysis of unevenly spaced paleoclimatic time series. *Computers & Geosciences*, 23(9): 929–945.
- Shukla, S. N.; and Marlin, B. 2018. Interpolation-Prediction Networks for Irregularly Sampled Time Series. In *Proceedings of International Conference on Learning Representations (ICLR)*.
- Shukla, S. N.; and Marlin, B. 2021. Multi-Time Attention Networks for Irregularly Sampled Time Series. In *Proceedings of International Conference on Learning Representations (ICLR)*.
- Shukla, S. N.; and Marlin, B. M. 2020. A survey on principles, models and methods for learning from irregularly sampled time series. *arXiv preprint arXiv:2012.00168*.
- Sun, C.; Li, H.; Song, M.; Cai, D.; Zhang, B.; and Hong, S. 2024. Time pattern reconstruction for classification of irregularly sampled time series. *Pattern Recognition*, 147: 110075.
- Tang, X.; Yao, H.; Sun, Y.; Aggarwal, C.; Mitra, P.; and Wang, S. 2020. Joint modeling of local and global temporal dynamics for multivariate time series forecasting with missing values. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 34, 5956–5963.
- Tashiro, Y.; Song, J.; Song, Y.; and Ermon, S. 2021. CSDI: Conditional Score-based Diffusion Models for Probabilistic Time Series Imputation. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, volume 34.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, volume 30.
- Wu, Y.; Ni, J.; Cheng, W.; Zong, B.; Song, D.; Chen, Z.; Liu, Y.; Zhang, X.; Chen, H.; and Davidson, S. B. 2021. Dynamic gaussian mixture based deep generative model for robust forecasting on sparse multivariate time series. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 35, 651–659.
- Wu, Z.; Pan, S.; Long, G.; Jiang, J.; Chang, X.; and Zhang, C. 2020. Connecting the dots: Multivariate time series forecasting with graph neural networks. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*, 753–763.
- Yalavarthi, V. K.; Madhusudhanan, K.; Scholz, R.; Ahmed, N.; Burchert, J.; Jawed, S.; Born, S.; and Schmidt-Thieme, L. 2024. GraFITi: Graphs for Forecasting Irregularly Sampled Time Series. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 38, 16255–16263.
- Yu, J.; Lin, Z.; Yang, J.; Shen, X.; Lu, X.; and Huang, T. S. 2018. Generative image inpainting with contextual attention. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5505–5514.
- Zhang, J.; Zheng, S.; Cao, W.; Bian, J.; and Li, J. 2023. Warformer: A Multi-scale Modeling Approach for Irregular Clinical Time Series. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*, 3273–3285.
- Zhang, W.; Yin, C.; Liu, H.; and Xiong, H. 2024a. Unleash the power of pre-trained language models for irregularly sampled time series. *arXiv preprint arXiv:2408.08328*.
- Zhang, W.; Yin, C.; Liu, H.; Zhou, X.; and Xiong, H. 2024b. Irregular Multivariate Time Series Forecasting: A Transformable Patching Graph Neural Networks Approach. In *Proceedings of International Conference on Machine Learning (ICML)*.
- Zhang, X.; Zeman, M.; Tsiligkaridis, T.; and Zitnik, M. 2021a. Graph-Guided Network for Irregularly Sampled Multivariate Time Series. In *Proceedings of International Conference on Learning Representations (ICLR)*.
- Zhang, Y.; Wang, X.; Yu, X.; Zhou, Z.; Xu, X.; Bai, L.; and Wang, Y. 2025. DiffODE: Neural ODE with Differentiable Hidden State for Irregular Time Series Analysis. In *2025 IEEE 41st International Conference on Data Engineering (ICDE)*, 2107–2120. IEEE Computer Society.
- Zhang, Z.-Y.; Zhang, S.-Q.; Jiang, Y.; and Zhou, Z.-H. 2021b. LIFE: Learning individual features for multivariate time series prediction with missing values. In *IEEE International Conference on Data Mining (ICDM)*, 1511–1516. IEEE.
- Zhou, T.; Niu, P.; Sun, L.; Jin, R.; et al. 2023. One fits all: Power general time series analysis by pretrained lm. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, 43322–43355.
- Zhou, Z.; Huang, Y.; Wang, Y.; Wu, Y.; Kwok, J.; and Liang, Y. 2025. Revitalizing Canonical Pre-Alignment for Irregular Multivariate Time Series Forecasting. *arXiv preprint arXiv:2508.01971*.