TIME-ADAPTIVE FRACTIONAL SPIKING NETWORK FOR IRREGULAR MULTIVARIATE TIME SERIES FORECASTING

Anonymous authors
Paper under double-blind review

ABSTRACT

Irregular multivariate time series (IMTS) forecasting is challenging due to irregular sampling, variable asynchronicity, and partial observability. Most existing methods for IMTS modeling primarily rely on temporal alignment or structural reformulation to alleviate the computational difficulties caused by irregular sampling. However, these methods often fail to fully exploit the sparse event characteristics arising from the partial observability of IMTS. With event-driven computation and sparse activation, spiking neural networks (SNNs) provide a natural modeling paradigm for preserving the event semantics of IMTS and capturing their temporal dependencies. Thus, we propose the Time-Adaptive Fractional Spiking Network (TAFNet), a multi-granularity framework for IMTS forecasting. We introduce a time-adaptive fractional leaky integrate-and-fire (TAF-LIF) neuron that jointly incorporates actual inter-observation intervals and a fractional memory kernel into the membrane potential evolution, enabling interval-aware temporal modeling and persistent historical memory. We further construct multi-granularity temporal representations to capture scale-dependent dynamics under varying observation densities, and integrate them with global periodic information extracted directly from irregular timestamps. Extensive experiments on real-world datasets demonstrate the state-of-the-art performance of TAFNet. The code is available at `https://anonymous.4open.science/r/TAFNet/`.

1 INTRODUCTION

Irregular multivariate time series (IMTS) are widely encountered in real-world applications, including healthcare (Sun et al., 2026), meteorology (Li et al., 2025), transportation (Wu et al., 2025a), and industrial sensing (Wang et al., 2026). Due to factors such as sensor failures, communication errors, and cost-driven acquisition policies, IMTS typically exhibit non-uniform sampling intervals within each variable, asynchronous observations across variables, and natural missingness. Moreover, valid observations are often distributed over only a small number of irregular timestamps, resulting in pronounced sparsity and non-uniformity along the continuous time axis. These inherent irregular sampling patterns and sparse observation distributions pose substantial challenges to IMTS forecasting.

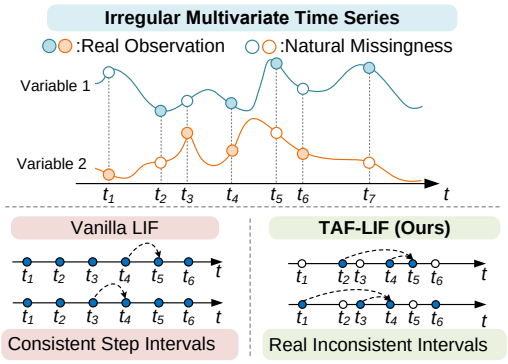

Figure 1: Illustration of vanilla LIF and TAF-LIF for irregular multivariate time series.

Recent studies have explored multiple directions to address the above challenges and model temporal dependencies under irregular observations. A common line of work maps irregular observations onto regular temporal grids through interpolation (Zhang et al., 2023; Zheng et al., 2026), temporal

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

alignment (Zhang et al., 2024a; Zhou et al., 2026b), or filling (Sun et al., 2026), thereby enabling standard sequence models to process the resulting representations. Another line of work employs neural ordinary differential equations (Kidger et al., 2020; Oh et al., 2024; Gao et al., 2026) or continuous-time state-space models (Brindle et al., 2025; Park et al., 2025) to characterize latent-state evolution in continuous time. In parallel, graph neural network models organize irregular observations as nodes, edges, or event sets to capture complex dependencies among asynchronous variables (Liu et al., 2025a; Luo et al., 2025). Although these strategies alleviate the computational difficulties caused by irregular sampling, they often fail to fully exploit the sparse and asynchronous event characteristics arising from the partial observability of IMTS. With event-driven computation and sparse activation, spiking neural networks (SNNs) can update neuronal states in response to valid observations. Therefore, SNNs provide a promising paradigm for IMTS forecasting.

Conventional leaky integrate-and-fire (LIF) neurons present two challenges for this purpose (Figure 1). First, a constant decay factor applies the same transition rule to observations separated by different temporal gaps. Second, their first-order recurrence summarizes the past in one membrane state, without an explicit mechanism for reweighting individual historical state increments according to elapsed time. Together, these properties motivate a neuronal update that couples interval adaptation with time-dependent historical weighting. Beyond neuronal dynamics, observation density and temporal patterns can vary substantially over time, motivating representations at multiple temporal resolutions. Hi-Patch already models multiple scales through hierarchical patch representations (Luo et al., 2025). In TAFNet, we organize observations into parallel temporal windows of different sizes, apply TAF-LIF encoding within these temporal contexts, and use global frequency features to guide their fusion. This organization connects multi-granularity aggregation with the proposed neuronal dynamics.

To address these limitations, we propose a multi-granularity time-adaptive fractional spiking network (TAFNet) for IMTS forecasting, centered on a novel time-adaptive fractional leaky integrate-and-fire (TAF-LIF) neuron. TAF-LIF combines real observation intervals with fractional history memory, allowing the membrane potential to preserve weighted contributions from historical events according to their temporal distances. Based on TAF-LIF, we design an event-aligned spike encoder that uses observation masks to constrain the input features, ensuring the TAF-LIF input contains valid observation events. Irregular observations are then organized at multiple temporal granularities and processed by patch embedding and mixing modules to capture local dependencies, inter-patch interactions, and cross-variable correlations. In parallel, a non-uniform discrete Fourier transform extracts global periodic patterns directly from irregular observations and fuses them with the multi-granularity representations. Finally, a lightweight decoder generates forecasts at future timestamps. Our contributions can be summarized as follows:

- We propose TAFNet, a time-adaptive fractional spiking framework for IMTS forecasting that explicitly models sparse observation events while capturing their temporal dependencies under irregular sampling.
- We introduce a novel TAF-LIF neuron that integrates real inter-observation intervals with fractional memory to capture interval-dependent temporal dependencies.
- We develop a multi-granularity temporal representation mechanism to capture scale-dependent dynamics under varying observation densities, and further integrate these representations using global frequency information.
- Extensive experiments on multiple real-world IMTS datasets demonstrate that TAFNet achieves better performance.

2 Related Work

IMTS Forecasting. Early IMTS studies focused on recurrent architectures such as GRU-D (Che et al., 2018), which employ time-decay mechanisms to account for missingness and irregular intervals. Continuous-time models, including Neural ODEs (Chen et al., 2018; Rubanova et al., 2019), further address irregular sampling by evolving latent states in continuous time. More recently, graph-based and attention-driven approaches have become prominent. Raindrop (Zhang et al., 2021) and ASTGI (Liu et al., 2025b) model asynchronous inter-variable dependencies via graph propagation, while tPatchGNN (Zhang et al., 2024a) models temporal patches and Hi-Patch (Luo et al., 2025) de-

velops a hierarchical patch representation across temporal scales. Frequency-domain methods, such as KAFNet (Zhou et al., 2026b) and TFMixer (Qiu et al., 2026), explore spectral representations through temporal alignment and learnable non-uniform Fourier transforms, respectively. Meanwhile, APN (Liu et al., 2026) introduces a time-aware soft aggregation strategy that alleviates interpolation errors in patch-based modeling.

SNNs for Temporal Modeling. Recent studies have proposed a series of spiking neural network architectures based on spike encoders for time series forecasting tasks, such as Spike-RNN, Spike-TCN, and iSpikformer (Lv et al., 2024). SpikF (Wu et al., 2025b) introduces frequency-domain analysis to enhance the capability of SNNs in capturing global dependencies. In addition, another line of research improves temporal representation by redesigning the intrinsic dynamics of spiking neurons. STC-LIF (Wang & Yu, 2024) and TS-LIF (Feng et al., 2025) introduce multi-timescale neuronal dynamics to capture temporal patterns at different ranges. LM-H (Hao et al., 2024) employs learnable state-transition parameters to regulate the contribution of historical states and current inputs, while multi-compartment neurons (Zheng et al., 2024) exploit dendritic heterogeneity to support multi-timescale information processing. More recently, f-LIF (Ge et al., 2025) formulates membrane evolution through fractional-order dynamics to enhance non-local temporal dependency modeling in spike sequences, whereas EA-LIF (Zhou et al., 2026a) develops event-synchronous neuronal updates for irregular telemetry forecasting. Despite these advances, existing spiking neurons do not fully characterize how irregular temporal distances affect historical-state accumulation during membrane-potential updates.

3 PRELIMINARIES

LIF neuron model. A standard neuron model is the LIF neuron, whose membrane potential accumulates input currents while decaying over time. Its discrete dynamics are formulated as

$$\bar{u}[t] = \beta u[t-1] + I[t], s[t] = H\left(\bar{u}[t] - V_{\mathrm{th}}\right), u[t] = \bar{u}[t] - s[t]V_{\mathrm{th}}, \tag{1}$$

where $I[t]$ denotes the input current, β is the membrane decay factor, V_{th} is the firing threshold, and $s[t] \in \{0, 1\}$ denotes the spike output.

Fractional Calculus. Fractional calculus extends classical differentiation and integration to non-integer orders, providing a natural mechanism for modeling history-dependent temporal dynamics. In this work, we adopt the Caputo fractional derivative. For $\alpha \in (0, 1)$, it is defined as

$$D_t^{\alpha} u(t) = \frac{1}{\Gamma(1-\alpha)} \int_0^t \frac{u'(s)}{(t-s)^{\alpha}}\, ds, \tag{2}$$

where $\Gamma(\cdot)$ denotes the gamma function. The kernel $(t-s)^{-\alpha}$ assigns distance-dependent weights to historical state changes, thereby introducing historical memory into the dynamics. Accordingly, replacing the first-order derivative in $\frac{dy(t)}{dt} = f(t, y(t))$ with the Caputo derivative yields

$$D_t^{\alpha} y(t) = f\big(t, y(t)\big). \tag{3}$$

4 METHODOLOGY

Figure 2 illustrates TAFNet, which combines time-adaptive fractional neuronal dynamics with multi-granularity temporal representations for IMTS forecasting. At its core, TAF-LIF couples interval-dependent membrane transitions with historical state increments weighted by their actual temporal distances. After masked normalization, the temporal branch partitions observations into multiple granularities and applies the spike encoder (EA-SE) to construct neuronal representations. These representations are processed by TTCN patch embedding and query-based patch mixing. In parallel, a NUDFT-based frequency branch extracts global spectral context to guide multi-granularity fusion. The NUDFT-based frequency module and query-based patch mixing follow TFMixer (Qiu et al., 2026). Finally, an MLP decoder maps the fused representation to forecasts, and masked denormalization restores the original data scale.

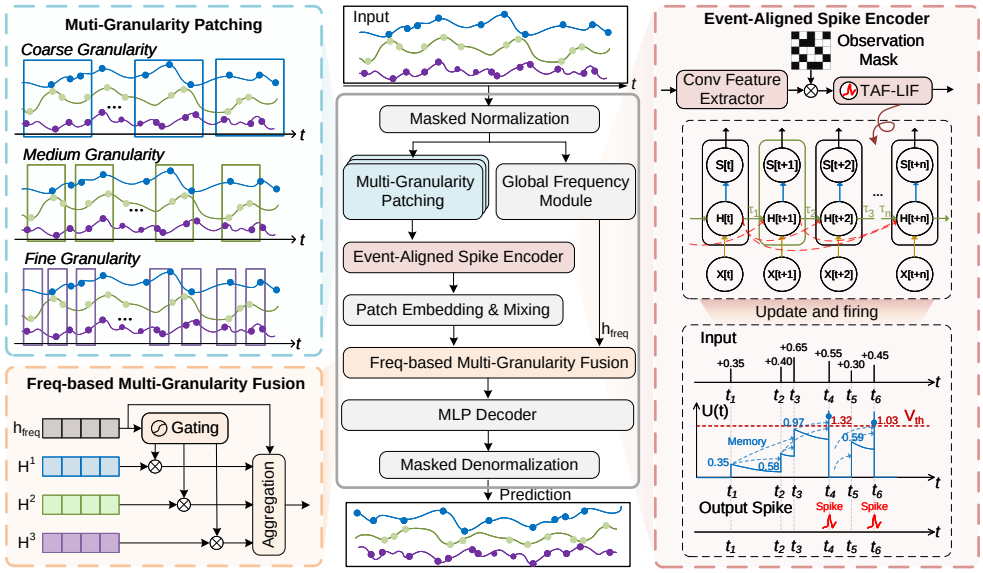

Figure 2: The main architecture of TAFNet. Irregular observations are organized into multiple temporal granularities, encoded by the observation-masked spike encoder, fused with global frequency features, and decoded for forecasting.

4.1 Global Frequency Module

We use a learnable Non-Uniform Discrete Fourier Transform (NUDFT) to extract global frequency features from irregular observations. Unlike the discrete Fourier transform, which requires uniformly sampled signals, NUDFT directly incorporates real timestamps into frequency analysis: $X(f_k) = \sum_{i=1}^{L} x_i e^{-j2\pi f_k t_i}$. This design avoids interpolation and preserves the intrinsic temporal structure of irregular sequences. The obtained frequency representation is then encoded by an MLP-based frequency encoder to generate global frequency features $h_{freq} = \mathrm{MLP}_{freq}(X(f))$.

4.2 Multi-Granularity Patching

To capture temporal patterns at different resolutions, we partition the irregular sequence into patches with multiple temporal granularities. Compared with a single patch size, multi-granularity partitioning preserves fine-grained local variations while capturing broader temporal structures. Given the encoded feature sequence $\mathbf{Z} = \{\mathbf{z}_i\}_{i=1}^{L}$ and the corresponding timestamps $\mathbf{T} = \{t_i\}_{i=1}^{L}$, we define a set of temporal window sizes $\mathcal{S} = \{s_1, s_2, \ldots, s_G\}$, where s_g denotes the window size of the g-th granularity. Let $t_{\min}$ and $t_{\max}$ denote the beginning and ending times of the observation interval. The number of patches at the g-th granularity is defined as

$$P_g = \left\lceil \frac{t_{\max} - t_{\min}}{s_g} \right\rceil . \tag{4}$$

For the p-th patch at the g-th granularity, we define its observation index set as

$$\mathcal{I}_p^{(g)} = \{ i \,|\, t_{\min} + (p-1)s_g \le t_i < t_{\min} + p s_g \} . \tag{5}$$

The corresponding patch is then represented as

$$\mathbf{P}_p^{(g)} = \left\{ \mathbf{z}_i \mid i \in \mathcal{I}_p^{(g)} \right\} . \tag{6}$$

4.3 Event-Aligned Spike Encoder

Time-Adaptive Fractional LIF. Vanilla LIF neurons are defined on fixed time intervals and use a constant membrane decay factor, which is unsuitable for IMTS with non-uniform observation gaps.

To address this limitation, we propose the TAF-LIF neuron. TAF-LIF explicitly incorporates real observation intervals into membrane potential evolution. Moreover, its fractional-order dynamics introduce history-dependent membrane updates, allowing the neuron to model temporal memory across historical observations. The dynamics of TAF-LIF are defined as

$$\tau_k D_t^\alpha u(t) = -u(t) + RI(t), \tag{7}$$

where R is the membrane resistance, τ_k is the membrane time factor, and D_t^α denotes the Caputo fractional derivative with order $\alpha \in (0, 1]$. When τ_k is constant and $\alpha = 1$, the fractional dynamics degenerate to the vanilla LIF dynamics.

For IMTS, observations of each variable are defined on a non-uniform temporal grid $\{t_0, t_1, \ldots, t_k\}$, where the interval $\Delta t_k = t_k - t_{k-1}$ varies over time. To adapt the neuronal dynamics to non-uniform observation intervals, the membrane time factor is computed as

$$\hat{\tau}_k = \hat{\tau}_0 + \mathrm{Softplus}\left(W\Delta t_k + b\right), \tag{8}$$

where $\hat{\tau}_0 > 0$ is a learnable parameter and $\tau_k = \hat{\tau}_k$ in the neuronal update. This formulation allows different observation gaps to induce different membrane-state transition behaviors.

To discretize Eq. equation 7 on such an irregular grid, we adopt the non-uniform L1 approximation of the Caputo derivative. We distinguish the charging membrane potential $\widetilde{u}_k$, before spike generation, from the post-reset membrane potential u_k. The fractional derivative is approximated as

$$D_t^\alpha u(t_k) \approx \; w_{k,k-1}^{(\alpha)}\left(\widetilde{u}_k - u_{k-1}\right) + \sum_{j=\ell_k}^{k-2} w_{k,j}^{(\alpha)}\left(u_{j+1} - u_j\right), \tag{9}$$

where $\ell_k = \max(0, k-K)$ and K counts the current transition together with at most $K-1$ historical increments. The corresponding non-uniform history coefficient is

$$w_{k,j}^{(\alpha)} = \frac{(t_k - t_j)^{1-\alpha} - (t_k - t_{j+1})^{1-\alpha}}{\Gamma(2-\alpha)(t_{j+1} - t_j)}. \tag{10}$$

Therefore, each historical membrane variation is weighted according to both its actual temporal distance from the current event and the fractional order. For notational convenience, we define the truncated fractional-memory term as

$$\mathcal{M}_k^{(\alpha)} = \sum_{j=\ell_k}^{k-2} w_{k,j}^{(\alpha)}\left(u_{j+1} - u_j\right). \tag{11}$$

Substituting Eq. equation 9 into Eq. equation 7 gives

$$\tau_k\left[w_{k,k-1}^{(\alpha)}\left(\widetilde{u}_k - u_{k-1}\right) + \mathcal{M}_k^{(\alpha)}\right] = -\widetilde{u}_k + RI_k. \tag{12}$$

Then, the TAF-LIF charging process is written as

$$\widetilde{u}_k = \frac{RI_k + \tau_k w_{k,k-1}^{(\alpha)} u_{k-1} - \tau_k \mathcal{M}_k^{(\alpha)}}{1 + \tau_k w_{k,k-1}^{(\alpha)}}. \tag{13}$$

In the reported experiments, α is learnable and initialized to 0.5. Equation equation 13 combines the preceding membrane state with the weighted historical term $\mathcal{M}_k^{(\alpha)}$. Its coefficients depend on the actual observation timestamps. With $K = 1$, the historical sum is empty; with $K = 2$, it contains at most one preceding state increment. Spike generation and reset follow Eq. equation 1, with $\widetilde{u}_k$ as the pre-reset potential.

Spiking Encoding Process. Given the p-th patch at the g-th temporal granularity $\mathbf{P}_p^{(g)}$, we first extract its local latent features using a convolutional layer followed by batch normalization:

$$\mathbf{F}_p^{(g)} = \mathrm{BN}\left(\mathrm{Conv}\left(\mathbf{P}_p^{(g)}\right)\right). \tag{14}$$

To set the neuronal input to zero at unobserved positions, the observation mask $\mathbf{M}_p^{(g)}$ is applied to the extracted features:

$$\widetilde{\mathbf{F}}_p^{(g)} = \mathbf{M}_p^{(g)} \odot \mathbf{F}_p^{(g)}, \tag{15}$$

where $\odot$ denotes element-wise multiplication. Subsequently, the masked features are used as the input current and the corresponding inter-observation intervals $\Delta\mathbf{T}_p^{(g)}$ are fed into the proposed TAF-LIF neuron to generate spike responses:

$$\widetilde{\mathbf{S}}_p^{(g)} = \text{TAF-LIF}\left(\widetilde{\mathbf{F}}_p^{(g)}, \Delta\mathbf{T}_p^{(g)}\right). \tag{16}$$

4.4 PATCH EMBEDDING AND MIXING

Since each patch remains an irregular subsequence with variable length and non-uniform internal intervals, we employ a Transformable Time-aware Convolution Network (TTCN) (Zhang et al., 2024b) to obtain its fixed-dimensional representation.

$$\mathbf{h}_p^{(g)} = \mathrm{TTCN}_g\left(\widetilde{\mathbf{S}}_p^{(g)}, \mathbf{T}_p^{(g)}\right). \tag{17}$$

The patch embeddings at the g-th granularity are collected as

$$\mathbf{H}^{(g)} = \left[\mathbf{h}_1^{(g)}, \mathbf{h}_2^{(g)}, \ldots, \mathbf{h}_{P_g}^{(g)}\right] \in \mathbb{R}^{P_g \times d}. \tag{18}$$

Finally, the multi-granularity patching representations are denoted by

$$\mathcal{H} = \left\{\mathbf{H}^{(1)}, \mathbf{H}^{(2)}, \ldots, \mathbf{H}^{(G)}\right\}. \tag{19}$$

Given the patch embeddings, we employ Patch Mixing Blocks to model temporal dependencies and variable interactions. Following the query-based patch mixing strategy of TFMixer (Qiu et al., 2026), Patch Attention first uses a set of learnable query tokens to adaptively attend to informative patches:

$$\mathbf{A}^{(g)} = \mathrm{Softmax}\left(\frac{\mathbf{Q}(\mathbf{H}^{(g)} + \mathbf{E}_{pos})^\top}{\sqrt{D}}\right), \quad \mathbf{U}^{(g)} = \mathbf{A}^{(g)}\mathbf{H}^{(g)}. \tag{20}$$

The obtained semantic patch tokens are then passed through a Patch Mixer and a Variable Mixer. The Patch Mixer applies an MLP along the patch dimension to model temporal interactions among semantic patches:

$$\mathbf{U}'^{(g)} = \mathrm{LayerNorm}\left(\mathbf{U}^{(g)} + \mathrm{MLP}_{patch}(\mathrm{Permute}(\mathbf{U}^{(g)}))\right). \tag{21}$$

Subsequently, the Variable Mixer applies another MLP along the variable dimension to exchange information across variables:

$$\mathbf{U}''^{(g)} = \mathrm{LayerNorm}\left(\mathbf{U}'^{(g)} + \mathrm{MLP}_{var}(\mathrm{Permute}(\mathbf{U}'^{(g)}))\right). \tag{22}$$

In this way, the Patch Mixing Blocks jointly capture informative temporal patch patterns and multivariate dependencies.

4.5 FREQUENCY-GUIDED MULTI-GRANULARITY FUSION (FMGF)

Since the number of patches $\{P_g\}_{g=1}^G$ may differ across granularities, each patch sequence $\mathbf{U}''^{(g)}$ is first processed by the corresponding patch mixing and aggregation module:

$$\mathbf{r}^{(g)} = \mathrm{Agg}_g\left(\mathbf{U}''^{(g)}\right), \qquad \mathbf{r}^{(g)} \in \mathbb{R}^d, \tag{23}$$

where $\mathbf{r}^{(g)}$ denotes the unified representation of the g-th temporal granularity. The representations from all granularities are stacked as

$$\mathbf{R} = \left[\mathbf{r}^{(1)}; \mathbf{r}^{(2)}; \ldots; \mathbf{r}^{(G)}\right] \in \mathbb{R}^{G \times d}. \tag{24}$$

Meanwhile, the global frequency representation $\mathbf{h}_{\mathrm{freq}}$ is used to generate granularity-specific gating weights:

$$\boldsymbol{\gamma} = \mathrm{Softmax}\left(\mathbf{W}_g\mathbf{h}_{\mathrm{freq}} + \mathbf{b}_g\right), \qquad \boldsymbol{\gamma} \in \mathbb{R}^G, \tag{25}$$

where γ_g indicates the importance of the g-th temporal granularity. The frequency-guided temporal representation is computed as

$$\mathbf{h}_{\mathrm{gate}} = \mathrm{Norm}\left(\sum_{g=1}^{G} \gamma_g \mathbf{r}^{(g)}\right). \tag{26}$$

This gating mechanism enables the model to emphasize temporal scales that are consistent with the global periodic characteristics of the input sequence. To preserve complementary information across all granularities, we further introduce a concatenation-based fusion branch:

$$\mathbf{h}_{\text{concat}} = \text{Norm}\left(\mathbf{W}_c\left[\mathbf{r}^{(1)} \| \mathbf{r}^{(2)} \| \cdots \| \mathbf{r}^{(G)}\right] + \mathbf{b}_c\right), \tag{27}$$

where $\|$ denotes feature concatenation. Finally, the multi-granularity temporal representation and the global frequency representation are integrated as

$$\mathbf{h}_{\text{fusion}} = \text{Norm}\left(\mathbf{h}_{\text{concat}} + \mathbf{h}_{\text{gate}} + \lambda \mathbf{h}_{\text{freq}}\right), \tag{28}$$

where λ is a learnable scaling parameter. Through this design, the proposed module adaptively selects informative temporal granularities according to global frequency characteristics while preserving complementary information across different temporal scales.

4.6 Forecasting Decoder

The Forecasting Decoder synthesizes the final prediction from the fused temporal and frequency representations. We employ an MLP-based decoder to project the learned representation to forecasts over the prediction horizon: $\hat{Y} = \text{MLP}(h_{\text{fusion}})$. Then, we use masked denormalization to map the prediction back to the original scale. Following Qiu et al. (2026), we jointly optimize the forecasting and reconstruction objectives:

$$\mathcal{L} = \mathcal{L}_{\text{fore}} + \mathcal{L}_{\text{recon}}, \tag{29}$$

where $\mathcal{L}_{\text{fore}}$ denotes the masked mean absolute error (MAE) computed over valid future targets, and $\mathcal{L}_{\text{recon}}$ denotes the masked MAE between the original historical observations and the historical sequence reconstructed by applying NUDFT followed by Inverse NUDFT.

5 Experiments

Datasets. We evaluate TAFNet on four widely used IMTS forecasting benchmarks, including PhysioNet, MIMIC, HumanActivity, and USHCN, covering healthcare, biomechanics, and climate domains. All datasets are split into 80% training, 10% validation, and 10% test.

Implementation Details. We evaluate TAFNet under the unified IMTS forecasting protocol of TFMixer (Qiu et al., 2026). All compared models use the same data splits, normalization procedure, and forecasting targets. ANN baseline results are taken from TFMixer, while SNN baseline results are obtained from our experiments under this same protocol. The lookback window lengths are set to 36 hours for PhysioNet and MIMIC, 3000 milliseconds for HumanActivity, and 3 years for USHCN. For the forecasting horizon, HumanActivity predicts the following 300 milliseconds, while the remaining datasets use the next three timestamps as forecasting targets. TAFNet results are reported as means and standard deviations over five random seeds. Further dataset and experimental implementation details are provided in Appendix A.

Baselines. We comprehensively compare our model with 15 representative baselines, including both recent IMTS forecasting methods and spike-based forecasting models. 1) *Advanced IMTS Forecasting Models.* We further compare with recent artificial neural network (ANN) based forecasting approaches, including NeuralFlows (Biloš et al., 2021), CRU (Schirmer et al., 2022), GNeuralFlow (Mercatali et al., 2024), tPatchGNN (Zhang et al., 2024a), GraFITi (Yalavarthi et al., 2024), Hi-Patch (Luo et al., 2025), KAFNet (Zhou et al., 2026b), APN (Liu et al., 2026), and TFMixer (Qiu et al., 2026). 2) *Spike-based Forecasting Models.* To evaluate the effectiveness of the proposed spiking framework, we additionally compare with representative spiking neural network baselines, including SpikeTCN, SpikeRNN, SpikeGRU, iSpikeformer (Lv et al., 2024), SpikeF (Wu et al., 2025b), and SEDformer (Zhou et al., 2026a).

Main Results. Table 1 reports forecasting results on four benchmarks. TAFNet achieves the lowest mean MAE on all four datasets and the lowest mean error on six of the eight dataset–metric pairs among the evaluated methods. Relative to the strongest reported ANN baseline for each metric, its mean MSE is 3.94% lower on USHCN and its mean MAE is 2.56% lower on MIMIC. TAFNet also attains lower mean MSE and MAE than all evaluated SNN baselines across the four datasets.

Dataset	HumanActivity		USHCN		PhysioNet		MIMIC	
Metric	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
NeuralFlows	0.1722±0.0090	0.3150±0.0094	0.2087±0.0258	0.3157±0.0187	0.4056±0.0033	0.4466±0.0027	0.6085±0.0101	0.5306±0.0066
CRU	0.1387±0.0073	0.2607±0.0092	0.2168±0.0162	0.3180±0.0248	0.6179±0.0045	0.5778±0.0031	0.5895±0.0092	0.5151±0.0048
GNeuralFlow	0.3936±0.1585	0.4541±0.0841	0.2205±0.0421	0.3286±0.0412	0.8207±0.0310	0.6759±0.0100	0.8957±0.0209	0.6450±0.0072
tPatchGNN	0.0443±0.0009	0.1247±0.0031	0.1885±0.0403	0.3084±0.0479	0.3133±0.0053	0.3697±0.0049	0.4431±0.0115	0.4077±0.0088
GraFITi	0.0437±0.0005	0.1221±0.0017	0.1691±0.0093	0.2777±0.0248	0.3075±0.0015	0.3637±0.0036	0.4359±0.0455	0.4142±0.0297
Hi-Patch	0.0435±0.0002	0.1204±0.0009	0.1749±0.0268	0.2717±0.0216	**0.3071±0.0029**	0.3675±0.0042	0.4279±0.0010	0.4033±0.0032
KAFNet	0.0429±0.0003	0.1161±0.0010	0.1698±0.0181	0.2690±0.0226	0.3164±0.0028	0.3715±0.0038	0.4402±0.0086	0.4102±0.0041
APN	0.0421±0.0001	0.1159±0.0006	0.1590±0.0137	0.2611±0.0167	0.3093±0.0011	0.3650±0.0026	0.4292±0.0027	0.4016±0.0016
TFMixer	0.0415±0.0004	0.1124±0.0005	0.1448±0.0044	0.2052±0.0036	0.3118±0.0023	0.3576±0.0019	**0.4182±0.0038**	0.3782±0.0034
SpikeTCN	0.0438±0.0006	0.1160±0.0007	0.2950±0.0024	0.3423±0.0007	0.3291±0.0012	0.3644±0.0006	0.6202±0.0029	0.4920±0.0017
SpikeGRU	0.0708±0.0001	0.1588±0.0001	0.2971±0.0013	0.3422±0.0002	0.4239±0.0002	0.4407±0.0001	0.6343±0.0009	0.4973±0.0002
SpikeRNN	0.0710±0.0001	0.1591±0.0001	0.2974±0.0029	0.3429±0.0017	0.3305±0.0018	0.3668±0.0011	0.6335±0.0007	0.4973±0.0001
iSpikeformer	0.0609±0.0002	0.1485±0.0003	0.2971±0.0032	0.3424±0.0006	0.3801±0.0393	0.4067±0.0313	0.6339±0.0011	0.4974±0.0002
SpikeF	0.0442±0.0005	0.1179±0.0006	0.1800±0.0136	0.2504±0.0083	0.3305±0.0022	0.3656±0.0014	0.5163±0.0022	0.4249±0.0012
SEDformer	0.0470±0.0005	0.1277±0.0018	0.1628±0.0096	0.2739±0.0075	0.3588±0.0016	0.4045±0.0011	0.5112±0.0086	0.4453±0.0054
TAFNet	**0.0413±0.0007**	**0.1119±0.0012**	**0.1391±0.0035**	**0.2020±0.0018**	0.3092±0.0011	**0.3545±0.0013**	0.4237±0.0041	**0.3685±0.0018**

Table 1: Forecasting performance on four IMTS datasets. Overall performance is evaluated by MSE and MAE (mean ± std). The best and second-best results are highlighted in **bold** and with an underline, respectively.

Dataset	HumanActivity		USHCN		PhysioNet		MIMIC	
Metric	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
w/o EA-SE	0.0435±0.0003	0.1127±0.0006	0.1484±0.0110	0.2067±0.0055	0.3102±0.0015	0.3552±0.0008	0.4297±0.0040	0.3694±0.0024
w/o TAF-LIF	0.0441±0.0007	0.1142±0.0013	0.1459±0.0028	0.2042±0.0023	0.3131±0.0029	0.3576±0.0025	0.4388±0.0036	0.3724±0.0030
w/o MGP	0.0438±0.0012	0.1136±0.0015	0.1401±0.0027	0.2044±0.0020	0.3117±0.0014	0.3576±0.0021	0.4382±0.0046	0.3710±0.0014
w/o FMGF	0.0435±0.0015	0.1125±0.0013	0.1416±0.0066	0.2027±0.0030	0.3112±0.0077	0.3570±0.0073	0.4348±0.0027	0.3711±0.0003
TAFNet	**0.0413±0.0007**	**0.1119±0.0012**	**0.1391±0.0035**	**0.2020±0.0018**	**0.3092±0.0011**	**0.3545±0.0013**	**0.4237±0.0041**	**0.3685±0.0018**

Table 2: Ablation studies for TAFNet. Best results are highlighted in **bold**.

Dataset	HumanActivity		USHCN		PhysioNet		MIMIC	
Metric	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
Vanilla LIF	0.0441±0.0007	0.1142±0.0013	0.1459±0.0028	0.2042±0.0023	0.3131±0.0029	0.3576±0.0025	0.4388±0.0036	0.3724±0.0030
f-LIF	0.0416±0.0006	0.1156±0.0024	0.1411±0.0069	0.2087±0.0143	0.3111±0.0030	0.3603±0.0044	0.4246±0.0039	0.3764±0.0017
EA-LIF	0.0418±0.0001	0.1163±0.0007	0.1442±0.0021	0.2081±0.0056	0.3112±0.0003	0.3608±0.0012	0.4326±0.0037	0.3771±0.0021
TAF-LIF	**0.0413±0.0007**	**0.1119±0.0012**	**0.1391±0.0035**	**0.2020±0.0018**	**0.3092±0.0011**	**0.3545±0.0013**	**0.4237±0.0041**	**0.3685±0.0018**

Table 3: Comparison of different spiking neurons in TAFNet. Best results are highlighted in **bold**.

Ablation Studies. Table 2 evaluates four architectural changes. The *w/o EA-SE* variant removes the entire spike encoder, including its TAF-LIF neuronal processing; it does not denote removal of the observation mask alone. This comparison evaluates the contribution of the complete encoder. The *w/o TAF-LIF* variant instead retains the encoder and replaces TAF-LIF with vanilla LIF. The full model obtains lower mean errors than both variants on all four datasets, supporting the role of the encoder and its neuronal design. The *w/o MGP* variant uses a single granularity, while *w/o FMGF* replaces frequency-guided multi-granularity fusion with direct feature addition.

Comparison of Spiking Neuron Models. Table 3 compares vanilla LIF, f-LIF, EA-LIF, and TAF-LIF within TAFNet. Each alternative neuron replaces TAF-LIF in the spike encoder. All other architectural components, data processing, training settings, and evaluation procedures remain unchanged across these variants. TAF-LIF achieves the lowest mean MSE and MAE across the four datasets in this comparison. Because the forecasting architecture is shared, these results provide a more focused assessment of the neuronal design than comparisons between different end-to-end networks. They support jointly incorporating interval adaptation and fractional historical weighting in the spike encoder.

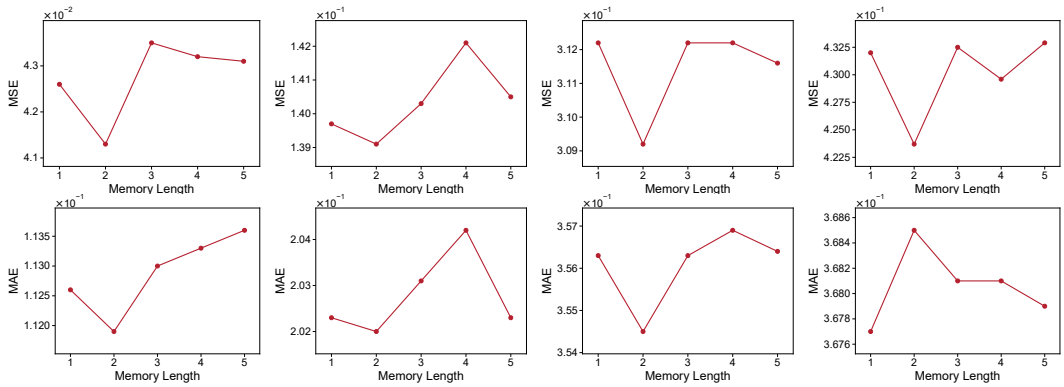

(a) HumanActivity (b) USHCN (c) PhysioNet (d) MIMIC

Figure 3: Parameter sensitivity studies of memory length in TAFNet. In each subfigure, the top and bottom panels show MSE and MAE, respectively.

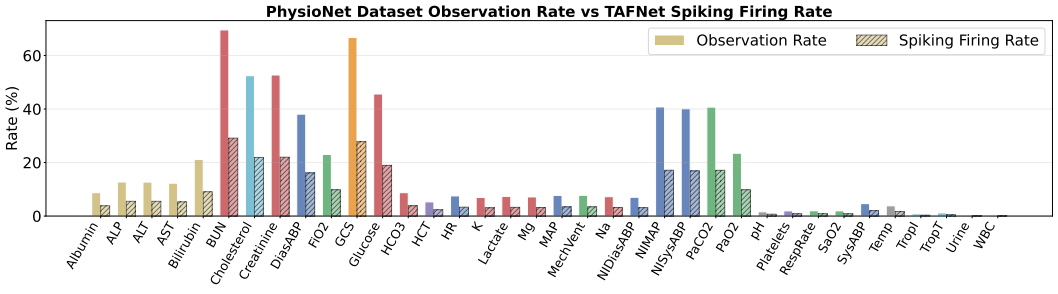

Figure 4: Comparison between the observation rate and the spike firing rate of TAFNet on each variable in the PhysioNet dataset. Solid bars represent the valid observation rate in the test set, while hatched bars denote the spike firing rate of TAFNet.

Parameter Sensitivity Analysis. Figure 3 examines the number of increments retained by TAF-LIF. The reported configuration uses $K = 2$ on all four datasets, corresponding to the current transition and at most one preceding increment. The results favor a compact history in this setting: extending the memory window does not yield a monotonic improvement. Here, K controls the number of retained increments, while the coefficients in Eq. equation 10 determine how their actual temporal distances influence the update.

Visualization. Figure 4 compares the valid observation rate and spike firing rate for each variable in PhysioNet. The firing rates broadly track differences in observation density across variables, illustrating the activity distribution of the spike encoder.

Due to space constraints, more experimental analysis and visualization are provided in Appendix B.

6 Conclusion

In this paper, we present TAFNet, a multi-granularity time-adaptive fractional spiking network framework specifically designed for IMTS forecasting. By integrating actual inter-observation intervals and fractional memory into the proposed TAF-LIF neuron, TAFNet effectively captures irregular temporal dynamics and local historical dependencies. Furthermore, the event-aligned spike encoder and multi-granularity patching mechanism enable the model to preserve sparse observation events and characterize temporal patterns under varying observation densities. Extensive experiments on multiple real-world IMTS datasets demonstrate that TAFNet achieves better performance.

AI USE STATEMENT

Generative AI tools assisted with language editing, manuscript restructuring, and the organization of revision priorities and proposed follow-up experiments. The experimental values reported in this manuscript were supplied by the authors; no new experimental results were generated during this text revision.

REFERENCES

Marin Biloš, Johanna Sommer, Syama Sundar Rangapuram, Tim Januschowski, and Stephan Günnemann. Neural flows: Efficient alternative to neural odes. In *Advances in Neural Information Processing Systems*, volume 34, pp. 21325–21337, 2021.

Benjamin Brindle, Thomas Derrick Hull, Matteo Malgaroli, and Nicolas Charon. Vista-ssm: Varying and irregular sampling time-series analysis via state-space models. *Psychological Methods*, August 2025. ISSN 1082-989X.

Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, David Sontag, and Yan Liu. Recurrent neural networks for multivariate time series with missing values. *Scientific reports*, 8(1):6085, 2018.

Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.

Shibo Feng, Wanjin Feng, Xingyu Gao, Peilin Zhao, and Zhiqi Shen. Ts-lif: A temporal segment spiking neuron network for time series forecasting. *arXiv preprint arXiv:2503.05108*, 2025.

Mengzhou Gao, Kaiwei Wang, and Pengfei Jiao. One-step graph-structured neural flows for irregular multivariate time series classification. *arXiv preprint arXiv:2605.10179*, 2026.

Chengjie Ge, Yufeng Peng, Zihao Li, Qiyu Kang, Xueyang Fu, Xuhao Li, Qixin Zhang, Junhao Ren, and Zheng-Jun Zha. Fractional-order spiking neural network. *arXiv preprint arXiv:2507.16937*, 2025.

Zecheng Hao, Xinyu Shi, Zihan Huang, Tong Bu, Zhaofei Yu, and Tiejun Huang. A progressive training framework for spiking neural networks with learnable multi-hierarchical model. In *International Conference on Learning Representations*, volume 2024, pp. 4925–4941, 2024.

Patrick Kidger, James Morrill, James Foster, and Terry Lyons. Neural controlled differential equations for irregular time series. In *Advances in Neural Information Processing Systems*, volume 33, pp. 6696–6707, 2020.

Yijun Li, Cheuk Hang Leung, and Qi Wu. Probabilistic learning of multivariate time series with temporal irregularity. *IEEE Transactions on Knowledge and Data Engineering*, 37(5):2874–2887, 2025.

Jiexi Liu, Meng Cao, and Songcan Chen. Timecheat: A channel harmony strategy for irregularly sampled multivariate time series analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 18861–18869, 2025a.

Xvyuan Liu, Xiangfei Qiu, Hanyin Cheng, Xingjian Wu, Chenjuan Guo, Bin Yang, and Jilin Hu. Astgi: Adaptive spatio-temporal graph interactions for irregular multivariate time series forecasting. *arXiv preprint arXiv:2509.23313*, 2025b.

Xvyuan Liu, Xiangfei Qiu, Xingjian Wu, Zhengyu Li, Chenjuan Guo, Jilin Hu, and Bin Yang. Rethinking irregular time series forecasting: A simple yet effective baseline. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 23873–23881, 2026.

Yicheng Luo, Bowen Zhang, Zhen Liu, and Qianli Ma. Hi-patch: Hierarchical patch gnn for irregular multivariate time series. In *Forty-second International Conference on Machine Learning*, 2025.

Changze Lv, Yansen Wang, Dongqi Han, Xiaoqing Zheng, Xuanjing Huang, and Dongsheng Li. Efficient and effective time-series forecasting with spiking neural networks. *arXiv preprint arXiv:2402.01533*, 2024.

Giangiacomo Mercatali, Andre Freitas, and Jie Chen. Graph neural flows for unveiling systemic interactions among irregularly sampled time series. In *Advances in Neural Information Processing Systems*, volume 37, pp. 57183–57206, 2024.

YongKyung Oh, Dongyoung Lim, and Sungil Kim. Stable neural stochastic differential equations in analyzing irregular time series data. In *The Twelfth International Conference on Learning Representations*, 2024.

Byoungwoo Park, Hyungi Lee, and Juho Lee. Amortized control of continuous state space feynman-kac model for irregular time series. In *International Conference on Learning Representations*, volume 2025, pp. 36795–36829, 2025.

Xiangfei Qiu, Kangjia Yan, Xvyuan Liu, Xingjian Wu, and Jilin Hu. Bridging time and frequency: A joint modeling framework for irregular multivariate time series forecasting. In *Forty-third International Conference on Machine Learning*, 2026.

Nitin Rathi and Kaushik Roy. Diet-snn: A low-latency spiking neural network with direct input encoding and leakage and threshold optimization. *IEEE Transactions on Neural Networks and Learning Systems*, 34(6):3174–3182, 2021.

Yulia Rubanova, Ricky TQ Chen, and David K Duvenaud. Latent ordinary differential equations for irregularly-sampled time series. *Advances in neural information processing systems*, 32, 2019.

Mona Schirmer, Mazin Eltayeb, Stefan Lessmann, and Maja Rudolph. Modeling irregular time series with continuous recurrent units. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 19388–19405, 2022.

Chenxi Sun, Moxian Song, Derun Cai, Baofeng Zhang, Hongyan Li, and Shenda Hong. A review of deep learning methods for irregularly sampled medical time series data. *Health Data Science*, 6:0456, 2026.

Lihao Wang and Zhaofei Yu. Autaptic synaptic circuit enhances spatio-temporal predictive learning of spiking neural networks. *arXiv preprint arXiv:2406.00405*, 2024.

Yakun Wang, Yisheng Zou, Gang Wang, Sonlin He, Linfu Sun, and Ning Zhou. A continuous-time generative model for irregular and regular industrial time series. *Pattern Recognition*, pp. 114429, 2026.

Fei Wu, Changjiang Zheng, Shiyu Zhou, Ye Lu, Zhilong Wu, and Shukang Zheng. An interpretable approach to passenger flow prediction and irregular passenger travel patterns understanding in metro system. *Expert Systems with Applications*, 265:125991, 2025a.

Wenjie Wu, Dexuan Huo, and Hong Chen. Spikf: Spiking fourier network for efficient long-term prediction. In *Forty-second International Conference on Machine Learning*, 2025b.

Vijaya Krishna Yalavarthi, Kiran Madhusudhanan, Randolf Scholz, Nourhan Ahmed, Johannes Burchert, Shayan Jawed, Stefan Born, and Lars Schmidt-Thieme. Grafiti: Graphs for forecasting irregularly sampled time series. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 16255–16263, 2024.

Man Yao, Jiakui Hu, Zhaokun Zhou, Li Yuan, Yonghong Tian, Bo Xu, and Guoqi Li. Spike-driven transformer. *Advances in neural information processing systems*, 36:64043–64058, 2023.

Weijia Zhang, Chenlong Yin, Hao Liu, Xiaofang Zhou, and Hui Xiong. Irregular multivariate time series forecasting: A transformable patching graph neural networks approach. In *Forty-first International Conference on Machine Learning*, 2024a.

Weijia Zhang, Le Zhang, Jindong Han, Hao Liu, Yanjie Fu, Jingbo Zhou, Yu Mei, and Hui Xiong. Irregular traffic time series forecasting based on asynchronous spatio-temporal graph convolutional networks. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 4302–4313, 2024b.

Xiang Zhang, Marko Zeman, Theodoros Tsiligkaridis, and Marinka Zitnik. Graph-guided network for irregularly sampled multivariate time series. *arXiv preprint arXiv:2110.05357*, 2021.

Xinlu Zhang, Shiyang Li, Zhiyu Chen, Xifeng Yan, and Linda Ruth Petzold. Improving medical predictions by irregular multimodal electronic health records modeling. In *International conference on machine learning*, pp. 41300–41313. PMLR, 2023.

Hanle Zheng, Zhong Zheng, Rui Hu, Bo Xiao, Yujie Wu, Fangwen Yu, Xue Liu, Guoqi Li, and Lei Deng. Temporal dendritic heterogeneity incorporated with spiking neural networks for learning multi-timescale dynamics. *Nature Communications*, 15(1):277, 2024.

Zinan Zheng, Yang Liu, and Jia Li. Mesh interpolation graph network for dynamic and spatially irregular global weather forecasting. *Advances in Neural Information Processing Systems*, 38: 25287–25311, 2026.

Ziyu Zhou, Yuchen Fang, Weilin Ruan, Shiyu Wang, James Kwok, and Yuxuan Liang. Sedformer: Event-synchronous spiking transformers for irregular telemetry time series forecasting. *arXiv preprint arXiv:2602.02230*, 2026a.

Ziyu Zhou, Yiming Huang, Yanyun Wang, Yuankai Wu, James Kwok, and Yuxuan Liang. Revitalizing canonical pre-alignment for irregular multivariate time series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 29115–29123, 2026b.

A Datasets, Data Preprocessing, and Implementation Details

A.1 Datasets

To ensure comprehensive and fair comparisons, we conduct experiments on four widely used benchmarks for irregular multivariate time series forecasting: PhysioNet, MIMIC-III, HumanActivity, and USHCN. These datasets cover diverse application domains, including healthcare, biomechanics, and climate science, and exhibit substantially different sampling densities, sequence lengths, and irregular observation patterns. PhysioNet Challenge 2012 contains multivariate clinical measurements collected during the first 48 hours of ICU stays, while MIMIC-III provides large-scale de-identified ICU records with sparse and heterogeneous observations. HumanActivity consists of three-dimensional positional measurements collected from subjects performing different physical activities, and USHCN contains long-term meteorological records from weather stations across the United States. For all datasets, we follow the preprocessing protocols adopted in previous IMTS forecasting studies (Qiu et al., 2026). All datasets are partitioned into training, validation, and test sets using an 80%, 10%, and 10% split, respectively. Detailed dataset statistics used in this paper are reported in Table 4.

Dataset	# Variables	# samples	Avg # Obs.	Max Length
PhysioNet	36	11,981	308.6	47
MIMIC	96	21,250	144.6	96
HumanActivity	12	1,359	362.2	131
USHCN	5	1,114	313.5	337

Table 4: Dataset statistics.

PhysioNet. The PhysioNet Challenge 2012 dataset contains 12,000 irregular multivariate time series, each corresponding to a unique ICU patient. For each patient, 41 clinical variables are recorded sporadically during the first 48 hours of ICU admission. Following prior studies, each sequence is divided into an observation window and a forecasting window.

MIMIC-III. The MIMIC-III dataset consists of de-identified clinical records from 23,457 ICU patients. Each patient is represented by 96 clinical variables collected during the first 48 hours of ICU admission. The dataset is characterized by sparse, asynchronous, and highly heterogeneous observations across variables.

HumanActivity. The HumanActivity dataset contains three-dimensional positional sensor measurements collected from five subjects performing different physical activities. It includes 12 variables observed at irregular timestamps. The continuous recordings are segmented into 5,400 samples, each covering a duration of 4,000 milliseconds.

USHCN. The United States Historical Climatology Network dataset contains more than 150 years of climate records collected from weather stations across the United States, with five meteorological variables. Following the standard preprocessing protocol, we use observations from 1,114 stations between 1996 and 2000, resulting in 26,736 irregular multivariate time-series instances.

A.2 Implementation Details

Following previous IMTS forecasting studies, we adopt Mean Squared Error (MSE) and Mean Absolute Error (MAE) as the evaluation metrics. The lookback horizons are set to 36 hours for PhysioNet and MIMIC-III, 3,000 milliseconds for HumanActivity, and 3 years for USHCN. For HumanActivity, the forecasting horizon is 300 milliseconds, while the next three timestamps are used as prediction targets for the remaining datasets, following the settings in prior works (Liu et al., 2026; Qiu et al., 2026). Our experiments for TAFNet and the SNN baselines are implemented in PyTorch 2.6.0 with CUDA 12.4 and conducted on an NVIDIA RTX 3090 GPU. TAFNet is optimized using the Adam optimizer. The forecasting term is computed using the masked MAE loss over valid future targets, while the reconstruction term constrains the NUDFT and Inverse NUDFT modules to reconstruct historical observations. TAFNet is evaluated with five random seeds, and its mean

Table 5: Hyperparameters for the proposed TAFNet model.

Hyperparameter	PhysioNet	MIMIC	HumanActivity	USHCN
Granularity	[2,4,8]	[2,4,6]	[4,6,8]	[4,6]
Model dimensions	128	64	64	64
Learning rate	0.01	0.001	0.01	0.001
Patch mixer layers	2	4	4	4
Variable mixer layers	1	1	1	1
Batch size	32	34	16	16
τ_0	2	2	2	2
Memory length K	2	2	2	2

Dataset	**HumanActivity**		**USHCN**		**PhysioNet**		**MIMIC**	
Metric	**MSE**	**MAE**	**MSE**	**MAE**	**MSE**	**MAE**	**MSE**	**MAE**
2	0.0438±0.0012	0.1136±0.0015	0.1401±0.0027	0.2044±0.0020	0.3117±0.0014	0.3576±0.0021	0.4382±0.0046	0.3710±0.0014
4	0.0444±0.0017	0.1147±0.0012	0.1418±0.0028	0.2041±0.0025	0.3216±0.0063	0.3661±0.0064	0.4438±0.0040	0.3732±0.0021
6	0.0443±0.0015	0.1148±0.0006	0.1411±0.0033	0.2032±0.0030	0.3226±0.0055	0.3678±0.0057	0.4448±0.0041	0.3752±0.0024
8	0.0441±0.0003	0.1149±0.0007	0.1414±0.0013	0.2036±0.0021	0.3227±0.0061	0.3674±0.0068	0.4484±0.0050	0.3765±0.0020
4 6	0.0436±0.0004	0.1131±0.0013	**0.1391±0.0035**	**0.2020±0.0018**	0.3159±0.0022	0.3606±0.0034	0.4388±0.0036	0.3724±0.0030
2 4 6	0.0422±0.0010	0.1127±0.0017	0.1410±0.0063	0.2036±0.0038	0.3133±0.0017	0.3563±0.0006	**0.4237±0.0041**	0.3685±0.0018
2 4 8	0.0427±0.0008	0.1122±0.0015	0.1423±0.0041	0.2045±0.0035	**0.3092±0.0011**	**0.3545±0.0013**	0.4334±0.0044	**0.3677±0.0010**
2 6 8	0.0437±0.0014	0.1127±0.0018	0.1502±0.0145	0.2048±0.0053	0.3168±0.0039	0.3616±0.0039	0.4314±0.0022	0.3680±0.0003
4 6 8	**0.0413±0.0007**	**0.1119±0.0012**	0.1468±0.0084	0.2055±0.0038	0.3166±0.0019	0.3595±0.0027	0.4372±0.0059	0.3707±0.0027

Table 6: Sensitivity analysis of different granularity configurations for TAFNet. Best results are highlighted in **bold**.

and standard deviation are reported. The ANN forecasting results, including their reported standard deviations, are taken from TFMixer (Qiu et al., 2026). The SNN forecasting results are obtained under the same data splits, normalization procedure, and forecasting targets.

B Additional Experiments and Analysis

B.1 Parameter Sensitivity Analysis

Table 6 investigates the effect of different temporal granularity settings. We compare single-granularity configurations, including $\{2\}$, $\{4\}$, $\{6\}$, and $\{8\}$, with multi-granularity configurations composed of different patch number combinations. All granularity configurations were selected exclusively according to the validation-set MSE. The test set was not involved in hyperparameter selection. The results show that multi-granularity modeling generally achieves better forecasting performance than using a single patch length. These improvements demonstrate that representations constructed at different temporal resolutions provide complementary information for capturing both local variations and broader temporal relations. We further evaluate several manually specified granularity combinations and observe that the preferred combination varies slightly across datasets, reflecting their different sampling patterns and temporal characteristics. Overall, the results consistently validate the effectiveness of the proposed multi-granularity representation mechanism over single-granularity modeling.

B.2 Comparison with Other Encoders

Table 7 compares EA-SE with GRU- and LSTM-based alternatives within the TAFNet forecasting backbone. EA-SE achieves the lowest mean MSE and MAE across all four datasets. Relative to the GRU variant, it reduces mean MSE by 4.40% on HumanActivity and 2.86% on USHCN. These results support the advantage of the TAF-LIF-based encoder over the evaluated conventional temporal encoders for irregular multivariate time series forecasting.

Dataset	HumanActivity		USHCN		PhysioNet		MIMIC	
Metric	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
GRU	0.0432±0.0009	0.1126±0.0012	0.1432±0.0047	0.2052±0.0037	0.3098±0.0014	0.3553±0.0019	0.4289±0.0053	0.3699±0.0018
LSTM	0.0435±0.0006	0.1131±0.0018	0.1433±0.0035	0.2060±0.0033	0.3099±0.0017	0.3549±0.0015	0.4291±0.0049	0.3702±0.0025
EA-SE	**0.0413±0.0007**	**0.1119±0.0012**	**0.1391±0.0035**	**0.2020±0.0018**	**0.3092±0.0011**	**0.3545±0.0013**	**0.4237±0.0041**	**0.3685±0.0018**

Table 7: Best results are highlighted in **bold**.

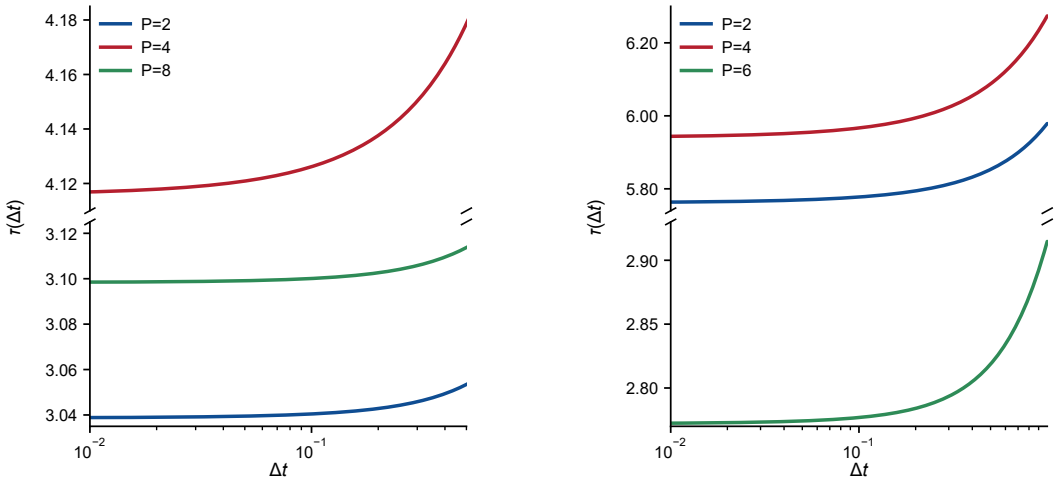

(a) PhysioNet (b) MIMIC

Figure 5: Relationship between the inter-observation interval and the time constant in TAF-LIF. Different curves correspond to different temporal granularities.

B.3 Visualization

To further examine the time-adaptive behavior of TAF-LIF, Fig. 5 visualizes the relationship between the inter-observation interval (Δt) and the adaptive time constant ($\tau(\Delta t)$) under different temporal granularities. The time constant changes continuously as the actual observation interval varies, demonstrating that TAF-LIF adjusts its membrane dynamics according to irregular temporal gaps rather than applying a fixed decay rule. Moreover, the distinct response curves associated with different granularity settings indicate that each temporal scale learns its own interval-dependent adaptation pattern. This enables TAF-LIF to characterize temporal dependencies at different resolutions while accounting for non-uniform observation intervals.

B.4 Computational Efficiency Analysis

We evaluate computational efficiency using the number of parameters, training time, FLOPs, SOPs, and estimated energy consumption per inference. A 32-bit floating-point operation and a synaptic accumulation are assumed to consume 4.6 pJ and 0.9 pJ, respectively, under a 45-nm CMOS process (Rathi & Roy, 2021; Yao et al., 2023).

FLOPs and SOPs characterize different computational paradigms. FLOPs represent the dense multiply-accumulate operations performed by ANN components, whereas SOPs represent the event-driven synaptic accumulations activated by spikes in SNN components. For the l-th spiking block, SOPs are estimated from its dense operation count, firing rate, and neuronal time steps as

$$\mathrm{SOPs}_l = f_{r,l} \times T_l \times \mathrm{FLOPs}_l^{\mathrm{SNN}}. \tag{30}$$

In our profiling implementation, the temporal dimension is already included in the operation count and is therefore not counted twice. Pure ANN models do not involve event-driven synaptic operations, and their SOP entries are thus marked as unavailable. For hybrid models, the FLOPs column

Method	Params (M)	Train Time (s/epoch)	FLOPs (G)	SOPs (G)	Energy (μJ/inference)
Hi-Patch	0.9960	50.471	6.648756	–	30584.277
TFMixer	0.1394	7.945	0.229392	–	1055.202
SEDformer	0.0213	13.371	0.010277	0.007756	54.256
TAFNet	0.3931	30.505	0.003166	0.091873	97.249

Table 8: Computational cost comparison on PhysioNet. Energy values are analytical estimates under the stated operation-cost model.

reports operations in the ANN components, while the SOPs column reports effective operations in the spiking components. Their inference energy is estimated by

$$E = 4.6\ \mathrm{pJ} \times \mathrm{FLOPs}_{\mathrm{ANN}} + 0.9\ \mathrm{pJ} \times \mathrm{SOPs}_{\mathrm{SNN}}. \tag{31}$$

Table 8 reports parameter counts, training time, operation counts, and an operation-based energy estimate on PhysioNet. TAFNet has 0.3931 million parameters and a reported training time of 30.505 seconds per epoch. The energy column is calculated under the stated per-operation cost model; it is an analytical estimate rather than a device-level energy measurement. Interpreting this estimate requires distinguishing dense operations in the ANN backbone from synaptic accumulations that can actually exploit binary spikes.