# HI-PATCH: HIERARCHICAL PATCH GNN FOR IRREG ULAR MULTIVARIATE TIME SERIES MODELING

Anonymous authors

Paper under double-blind review

#### Abstract

Multi-scale information is crucial for multivariate time series modeling. However, most existing time series multi-scale analysis methods treat all variables in the same manner, which is not well adaptive to Irregular Multivariate Time Series (IMTS), where different variables have distinct original scales/sampling rates. Therefore, extracting temporal and inter-variable dependencies at multiple scales in IMTS remains challenging. To fill this gap, we propose a hierarchical patch graph network Hi-Patch. The key components of Hi-Patch are an intra-patch graph layer and several inter-patch graph layers. The intra-patch graph layer flexibly represents and fully captures both the local temporal and inter-variable dependencies of densely sampled variables at the original scale by employing fully connected graph networks within each patch, and obtains patch-level node representations through aggregation. Subsequently, several inter-patch graph layers are stacked to form a hierarchical architecture, where each layer updates specific patch-level nodes through scale-specific graph networks, progressively representing and extracting more global temporal and inter-variable features of both sparsely and densely sampled variables, and further aggregating to produce the next patch-level node representations. The output of the last inter-patch graph layer is fed into task-specific decoders to adapt to different downstream tasks. Experimental results on 8 datasets show that Hi-Patch outperforms a range of state-of-the-art models in both IMTS forecasting and classification tasks.

029 030 031

032

004

010 011

012

013

014

015

016

017

018

019

021

024

025

026

027

028

## 1 INTRODUCTION

Time series analysis has important applications in various fields such as healthcare (Lan et al., 2024), climate forecasting (Verma et al., 2024) and traffic planning (Li et al.). Due to the complexity and non-stationarity of real-world systems (Wang et al., 2024a), multivariate time series often exhibit different variations and fluctuations at different scales (Chen et al., 2024a). Previous studies (Wang et al., 2024b; Chen et al., 2024a; Cai et al., 2024) have demonstrated that effectively capturing multi-scale features of multivariate time series is essential for time series modeling.

Due to sensor malfunctions, varying sampling sources, or human factors, the final sampled time series often exhibits irregularities in real-world applications, resulting in Irregular Multivariate Time Series (IMTS). The peculiarity of IMTS lies in two aspects: firstly, the sampling interval within each variable is uneven. Secondly, asynchronous observations of different variables have distinct sampling rates. The subgraph in the bottom left corner of Figure 1 illustrates an IMTS example with two variables.

As a unique type of multivariate time series, IMTS also exhibits the inherent multi-scale characteristics of time series. Unfortunately, many of the existing multi-scale analysis methods have limitations in mining the multi-scale characteristics of IMTS. The mainstream approach of these methods is to downsample or segment the original series based on several fixed time steps (Challu et al., 2023; Shabani et al., 2023), thereby forming several views that reflect features at different scale levels, as shown in the top row of Figure 1. On this basis, global temporal and inter-variable dependencies are extracted from coarse-grained views, while local dependencies are extracted from fine-grained views. This approach assumes that all variables have a consistent original scale so that each downsampled view consistently reflects the features of all variables at the same granularity level, thereby achieving consistent separation of multi-scale patterns for all variables.



Figure 1: Comparison of downsampling results of two variables for a regular (top row) and an irregular (bottom row) multivariate time series with step sizes of 1 (left column), 2 (middle column) and 4 (right column). Regular time series exhibit a consistent feature level at a specific downsampling scale, while irregular time series contain mixed feature levels at a specific downsampling scale.

However, in IMTS, some variables have only a few sampling points and do not possess fine-grained 073 features themselves, making their downsampled view meaningless (e.g., the variable  $V_0$  in Figure 074 1). Some other variables are densely sampled and require multiple downsampling steps to reflect 075 global patterns (e.g., the variable  $V_1$  in Figure 1). In such cases, it is difficult to determine a con-076 sistently suitable downsampling level for all variables, making existing methods infeasible from the 077 beginning. Furthermore, in IMTS, a downsampled view can contain mixed granularity features, as shown in the bottom row of Figure 1, but existing methods tend to extract and analyze a consistent 079 feature at a specific scale, which are difficult to handle mixed features. Additionally, despite the distinct scales/sampling rates occurring among variables in IMTS, they are not entirely independent but 081 exhibit correlations(Zhang et al., 2022). Specifically, variables in MTS may display different intervariable correlations at different time scales (Cai et al., 2024). How to extract multi-scale temporal 083 dependencies while considering inter-variable correlations across different scales in IMTS, where variables have distinct original scales, remains a major challenge in IMTS modeling. 084

085 To address this gap, we propose Hi-Patch, a hierarchical patch graph network. Hi-Patch is based on the concept of patching, which has been recently proven to be effective in capturing local tem-087 poral dependencies (Nie et al., 2023; Chen et al., 2024a; Zhang et al., 2024). Hi-Patch first divides 880 the original observation set into multiple patches based on a short time span. Then, an intra-patch graph layer flexibly represents and fully captures local temporal dependencies of densely sampled variables and their correlations with other variables at the original scales by employing a fully con-090 nected graph network within each patch. Subsequently, to capture dependencies at a larger scale, 091 the nodes are aggregated to obtain patch-level feature nodes as inputs for several inter-patch graph 092 layers. Each inter-patch graph layer receives specific patch-level nodes which are obtained through node aggregation of the previous layer, thereby forming a hierarchical architecture. These inter-094 patch graph layers gradually extract more global temporal and inter-variable dependencies that both densely and sparsely sampled variables possess through scale-specific graph networks. The output 096 of the last inter-patch graph layer is finally fed into specific task decoders to adapt to different downstream tasks. During this process, each intra/inter-patch graph layer handles features at a specific 098 scale. Fine-grained features of densely sampled variables are extracted in the lower layers, while 099 coarse-grained features of all variables are extracted in the upper layers, achieving complete extraction of mixed multiple granularities temporal and inter-variable correlations in IMTS. Our main 100 contributions are summarized as follows: 101

102 103

068

069

- We introduce the intra-patch/inter-patch graph layers to flexibly represent and fully extract features of specific variables in IMTS at specific scales.
- We propose Hi-Patch, using hierarchical architecture to effectively achieve multi-scale modeling of IMTS from fine-grained to coarse-grained.
- We conduct experiments on IMTS forecasting and classification on 8 datasets, and the evaluation results demonstrate that Hi-Patch outperforms existing methods in most cases.

## 108 2 RELATED WORK

110 Irregular Multivariate Time Series Modeling Existing methods can be broadly categorized into 111 interpolation-based and raw-data-based methods. The former, employing methods such as kernel-112 based approaches (Shukla & Marlin, 2019; Wu et al., 2021), Gaussian process (Tan et al., 2021) or 113 hourly aggregation (Ma et al., 2020), aims to obtain a set of regularly spaced observations. How-114 ever, interpolation may lead to losing valuable information from the original series. Raw-data-based methods learn directly from IMTS. (Che et al., 2018) improve recurrent neural networks to accom-115 modate uneven time intervals, while (Horn et al., 2020; Shukla & Marlin, 2021) introduce time 116 embeddings to represent arbitrary sampling timestamps. (Rubanova et al., 2019; De Brouwer et al., 117 2019; Biloš et al., 2021; Schirmer et al., 2022; Chen et al., 2024b) use neural ordinary differential 118 equations to handle irregularities. Recent research has also introduced attention (Jhin et al., 2021) or 119 graph neural networks (Zhang et al., 2022; Yalavarthi et al., 2024; Zhang et al., 2024) to account for 120 inter-variate correlations in IMTS. Despite these advancements, most of these methods focus only 121 on single-scale characteristics of IMTS, and how to fully capture multi-scale features within IMTS 122 remains a challenge.

123

124 Multi-scale Modeling for Time Series Multi-scale information has proven to be crucial for time 125 series modeling. TAMS-RNNs (Chen et al., 2021) use different frequencies to update several sub-126 hidden states of RNNs to capture multi-scale information. Pyraformer (Liu et al., 2022) introduces 127 pyramid attention to extract features at varying temporal scales. N-HiTS (Challu et al., 2023) employs multi-rate data sampling and hierarchical interpolation to extract features at different reso-128 lutions. Scaleformer (Shabani et al., 2023) allocates forecasting models across different temporal 129 scales. Pathformer (Chen et al., 2024a) proposes a multi-scale Transformer with adaptive paths, 130 while TimeMixer (Wang et al., 2024a) achieves complementary predictive capabilities by disentan-131 gling variation at multiple scales. However, these methods are designed for regular MTS and are 132 not adaptive for IMTS. Although Warpformer (Zhang et al., 2023) presents a multi-scale approach 133 for IMTS, its original goal is to balance the differences in sampling densities across variables and it 134 involves upsampling interpolation for sparse variables, which may distort the original data distribu-135 tion. 136

Graph Neural Networks for Multivariate Time Series A series of studies have integrated GNN with various time series modeling frameworks to effectively capture inter-variable dependencies in MTS (Jin et al., 2024). These approaches have been widely applied in diverse domains, includ-ing transportation (Rahmani et al., 2023), healthcare (Wang et al., 2022), economics (Wang et al., 2021), demonstrating promising results. However, many of these methods are primarily designed for modeling synchronous correlations among variables, lacking sufficient extraction of asynchronous dependencies widely present in IMTS.

144 145

146

## **3 PROBLEM DEFINITION**

**147 Definition 1 (Irregular Multivariate Time Series)** We consider a dataset  $\mathcal{D}$  consisting of n IMTS **148** sample. Each sample of  $\mathcal{D}$  is a tuple, i.e.,  $\mathcal{D} := \{(\mathcal{S}_1, y_1), \dots, (\mathcal{S}_n, y_n)\}$ , where  $\mathcal{S}_i$  denotes the *i*-th time series sample and  $y_i \in \{1, \dots, C\}$  is the class label (C is the number of categories). We describe *i*-th sample  $\mathcal{S}_i$  as a set of  $M = |\mathcal{S}_i|$  observations such that  $\mathcal{S}_i := \{o_1, \dots, o_M\}$ . Each observation  $o_j$  is a tuple  $(t_j, z_j, v_j)$ , consisting of a timestamp  $t_j \in \mathbb{R}^+$ , an observed value  $z_j \in \mathbb{R}$ and a variable indicator  $v_j \in \{1, \dots, V\}$ , where V represents the total number of variables. An IMTS sample can thus be represented as:

154

156

157

158

$$S_i := \{ (t_j, z_j, v_j) | j = 1, ..., M \},$$
(1)

**Problem 1 (Irregular Multivariate Time Series Classification).** Given an IMTS sample  $S_i$ , the problem is to correctly predict its class label  $y_i$ :

$$\mathcal{C}(\mathcal{S}_i) \to y_i,$$
 (2)

where  $C(\cdot)$  denotes the classification model we aim to learn.

161 **Problem 2 (Irregular Multivariate Time Series Forecasting).** Given a split timestamp  $t_S$ , each sample  $S_i$  is segmented into a historical window  $\mathcal{X}_i := \{(t_j, z_j, v_j) | j = 1, ..., M, t_j \leq t_S\}$  and

a forecasting window  $\mathcal{Y}_i := \{[(t_j, v_j), z_j] | j = 1, ..., M, t_j > t_S\}$ . Elements  $t_j$  and  $v_j$  of j-th observation tuple in the set of forecasting windows are combined into a forecasting query  $q_j \in \mathcal{Q}_i$ . The problem is to accurately predict the corresponding observation values  $\mathcal{Z}_i$  correspondence to forecasting query  $\mathcal{Q}_i$  based on the historical window  $\mathcal{X}_i$ :

$$F(\mathcal{X}_i, \mathcal{Q}_i) \to \mathcal{Z}_i,$$
(3)

where  $\mathcal{F}(\cdot)$  denotes the forecasting model we aim to learn.



Figure 2: The model framework of Hi-Patch. The example input is an IMTS with three variables and a historical window of 0-40s which is divided into N = 4 patches with a patch size P = 10s. First, each observation is encoded into a graph node through an observation encoder. Then, fine-grained features of densely sampled variables within each patch are extracted via an intra-patch graph layer and nodes are aggregated into patch-level nodes as inputs for multiple stacked inter-patch graph layers. After passing through  $\lceil \log_2 N \rceil = 2$  inter-patch graph layers, coarse-grained features of all variables are progressively extracted and a single node embedding for each variable is obtained. Finally, the task decoder computes the downstream task output based on these node embeddings.

193 194 195

196

202

203

186

187

188

189

190

191

192

166 167

168

#### 4 Methodology

In this section, we will take the *i*-th sample  $S_i$  with M historical observations as an example for introducing our method. We introduce the observation encoder in Section 4.1, followed by the introduction of the single intra-patch graph layer and inter-patch graph layer in Sections 4.2 and 4.3, respectively. Section 4.4 describes how the hierarchical architecture is implemented, and Section 4.5 covers the task decoder. Figure 2 presents the overall architecture of our model.

#### 4.1 Observation Encoder

First, we introduce how to encode each observation into an  $d_{\text{model}}$ -dimensional graph node embedding. The *j*-th historical observation  $o_j$  of  $S_i$  corresponds to the tuple  $(t_j, v_j, z_j)$ , representing the observation value  $z_j$  of variable  $v_j$  at timestamp  $t_j$ . We encode these three parts separately.

For time encoding, we adopt continuous time embedding (Shukla & Marlin, 2021), which is designed specifically for IMTS. The *d*-th dimension of  $d_{\text{model}}$ -dimensional time embedding  $\phi(t_j)$  for  $t_j$  is calculated by:

$$\phi(t_j)[d] = \begin{cases} \omega_0 \cdot t + \alpha_0, & \text{if } d = 0\\ \sin(\omega_d \cdot t + \alpha_d), & \text{if } 0 < d < d_{\text{model}} \end{cases}, \tag{4}$$

213 where  $\omega_0$ ,  $\alpha_0$ ,  $\omega_d$  and  $\alpha_d$  are learnable parameters. For variable encoding, we randomly ini-214 tialize a learnable variable embedding for each variable, forming a variable embedding matrix 215  $E \in \mathbb{R}^{V \times d_{\text{model}}}$ , the embedding corresponding to variable  $v_j$  is  $E_{v_j}$ . For value encoding, a linear 216 layer  $f(\cdot)$  is used to map the observation value  $z_j$  into  $d_{\text{model}}$ -dimensional embedding. With these three embeddings, the graph node embedding for *j*-th observation  $o_j$  is calculated as:

$$h_j = \sigma_1[\phi(t_j) + \boldsymbol{E}_{v_j} + f(z_j)] \in \mathbb{R}^{d_{\text{model}}},\tag{5}$$

where  $\sigma_1(\cdot)$  is ReLU activation function.

#### 4.2 INTRA-PATCH GRAPH LAYER

In this section, we will introduce how to extract fine-grained features of densely sampled variables through an intra-patch graph layer. First, we divide the historical observation nodes into several non-overlapping patches based on a short time span P. Given a total historical time span of T, it can be divided into  $N = \begin{bmatrix} T \\ P \end{bmatrix}$  patches. The *n*-th patch includes observation nodes of all variables within the period from  $(n-1) \cdot P$  to  $n \cdot P$ , the index set of observation nodes in *n*-th patch is:

$$\mathcal{I}_n = \{ j \mid j = 1, 2, ..., M, (n-1) \cdot P < t_j \le n \cdot P \}.$$
(6)

The corresponding set of node states is  $\mathcal{H}_n = \{h_j \mid j \in \mathcal{I}_n\}$ , initialized by node embeddings. Each patch, divided based on a short time span, primarily consists of observation nodes from densely sampled variables, with a few or no observation nodes from sparsely sampled variables, thus forming a fine-grained view of the densely sampled variables.

#### 4.2.1 UPDATE

218

219 220

221

232 233

234

235

236

237

238

239

244

248

249

263 264 After obtaining several patches as fine-grained views of the densely sampled variables, each pair of nodes within a patch is connected by an edge to form a fully connected graph (intra-patch graph) for flexible representation of three types of local dependencies in IMTS, namely: 1) the same variable at different times (SVDT), 2) different variables at the same time (DVST) and 3) different variables at different times (DVDT). The set of edges in *n*-th patch is:

$$\mathcal{E}_n = \{ (u, v) \mid u \in \mathcal{I}_n, v \in \mathcal{I}_n \}.$$
(7)

We use graph attention network (GAT) (Veličković et al., 2018) to update the node states within each intra-patch graph for sufficiently extracting these dependencies. The formula for updating the state of node v is:

$$h_v = h_v + \sum_{u \in \mathcal{N}(v)} \mathbf{MHA}(h_v, h_u, h_u), \tag{8}$$

where  $\mathcal{N}(v)$  denotes the set of neighboring nodes of node v and **MHA** denotes the multi-head attention mechanism (Vaswani et al., 2017):

$$\mathbf{MHA}(h_v, h_u, h_u) = Softmax(\frac{(W_k h_u)^T (W_q h_v)}{\sqrt{d_{\text{model}}}}) \cdot W_v h_u, \tag{9}$$

Due to the existence of three different types of dependencies, we adopted three sets of  $\{W_q, W_k, W_v\}$  parameters to handle them separately. Compared to conventional methods that separately extract features along the time and variable dimensions, our intra-patch graph more comprehensively accounts for the asynchronous characteristics of IMTS. After the *L*-layer state update with Eq.(8), the local dependencies of the densely sampled variables at the original scale are fully extracted.

#### 255 256 4.2.2 AGGREGATION

To further extract features at a larger scale, we aggregate nodes of the same variables within each patch to create the overall feature nodes of variables during the period of each patch. Specifically, for the nodes of the *v*-th variable in the *n*-th patch, the index set of nodes is  $\mathcal{I}_{n,v} = \{j \mid j \in \mathcal{I}_n, v_j = v\}$ . We first calculate the average observation timestamp as a reference time  $\overline{t}_{n,v}^{(0)} = \frac{1}{|\mathcal{I}_{n,v}|} \sum_{j \in \mathcal{I}_{n,v}} t_j$ , and then use multi-time attention (Shukla & Marlin, 2021) to aggregate the nodes to  $\overline{t}_{n,v}^{(0)}$ :

$$h_{n,v}^{(0)} = \sum_{j \in \mathcal{I}_{n,v}} \mathbf{MHA}(\phi(t_j), \phi(\bar{t}_{n,v}^{(0)}), h_j).$$
(10)

When a variable has no observation within a patch, no aggregation node is created, thus preserving the irregularities of the original series. We use a mask indicator  $m_{n,v}^{(0)}$  to indicate whether the aggregation node exists of the v-th variable at the n-th patch. If  $|\mathcal{I}_{n,v}| = 0$ ,  $m_{n,v}^{(0)}$  is set as 0, else  $m_{n,v}^{(0)} = 1$ . The new set of node states for all N patches of all V variables after aggregation is  $\mathcal{H}^{(0)} = \{h_{n,v}^{(0)} \mid n = 1, 2, ..., N, v = 1, 2, ..., V, m_{n,v}^{(0)} = 1\}$ .

# 4.3 INTER-PATCH GRAPH LAYER271

After passing through an intra-patch graph layer, fine-grained features with time scales less than Pof the original series have been extracted. In this section, we aim to extract features with time scale P based on the aggregated node set  $\mathcal{H}^{(0)}$  through an inter-patch graph layer.

#### 4.3.1 UPDATE

275

276

289

300

301 302

303 304

305

306

307 308

Each node in the set  $\mathcal{H}^{(0)}$  represents the overall feature within time span P, and the reference time interval between nodes in adjacent patches is approximately equal to P. We connect the nodes located in adjacent patches pairwise in node set  $\mathcal{H}^{(0)}$  by an edge to form a P-scale graph (inter-patch graph). This inter-patch graph flexibly represents the temporal and inter-variable dependencies of densely sampled variables and some sparsely sampled variables with a sampling interval less than P at scale P. We continue to update the node states in this P-scale graph using GAT to extract the dependencies at scale P. The formula for updating is:

$$h_{n,v}^{(0)} = h_{n,v}^{(0)} + \sum_{n' \in \{n-1,n+1\}} \sum_{v' \in \{1,2,\dots,V\}} \mathbf{MHA}(h_{n,v}^{(0)}, h_{n',v'}^{(0)}, h_{n',v'}^{(0)}),$$
(11)

## 4.3.2 Aggregation

After the node state update, the features at scale P are extracted. Subsequently, to extract features at scale 2P, we aggregate every two nodes located on adjacent patches of the same variable by multitime attention again to obtain the input of the next inter-patch graph layer. Specifically, we calculate the average timestamp between two adjacent nodes as the reference time  $\bar{t}_{n,v}^{(1)} = \frac{\bar{t}_{2n-1,v}^{(0)} + \bar{t}_{2n,v}^{(0)}}{m_{2n-1,v}^{(0)} + m_{2n,v}^{(0)}}$ , where n = 1, 2, ..., N/2, and then we aggregate the adjacent nodes to this reference time through multi-time attention to obtain feature nodes of 2P scale:

$$h_{n,v}^{(1)} = \sum_{j \in \{2n-1,2n\}} \mathbf{MHA}(\phi(\bar{t}_{j,v}^{(0)}), \phi(\bar{t}_{n,v}^{(1)}), h_{j,v}^{(0)}).$$
(12)

The node state set after aggregating is  $\mathcal{H}^{(1)} = \{h_{n,v}^{(1)} \mid n = 1, 2, ..., N/2, v = 1, 2, ..., V, m_{n,v}^{(1)} = 1\},\$ which is served as input for the next inter-patch graph layer (2*P*-scale).

#### 4.4 HIERARCHICAL ARCHITECTURE

Section 4.3 describes the first inter-patch graph layer (*P*-scale). This layer takes  $\mathcal{H}^{(0)}$  as input, updates the node states, and aggregates them to  $\mathcal{H}^{(1)}$  as the input of next inter-patch graph layer (2*P* scale). Generalizing to the general case, the formula for the *l*-th inter-patch graph layer is:

$$\mathcal{H}^{(l)} = \text{Aggregation}(\text{Update}(\mathcal{H}^{(l-1)})), \tag{13}$$

where  $\mathcal{H}^{(l)} = \{h_{n,v}^{(l)} \mid n = 1, 2, ..., N/2^l, v = 1, 2, ..., V, m_{n,v}^{(l)} = 1\}$  and  $\mathcal{H}^{(l-1)} = \{h_{n,v}^{(l-1)} \mid n = 1, 2, ..., N/2^{l-1}, v = 1, 2, ..., V, m_{n,v}^{(l-1)} = 1\}$ . The recursive form of Eq.(13) allows us to build 309 310 311 a hierarchical architecture by stacking inter-patch graph layers, thereby further extracting features 312 at scales of 2P, 4P, 8P, 16P, .... Until  $l = \lceil \log_2 N \rceil$ , the output of this layer is  $\mathcal{H}^{(\lceil \log_2 N \rceil)} =$ 313  $\{h_{n,v}^{(\lceil \log_2 N \rceil)} \mid n = 1, v = 1, 2, ..., V, m_{n,v}^{(\lceil \log_2 N \rceil)} = 1\}$ , where a single node embedding is obtained 314 for each variable. This process involves a total of  $\lceil \log_2 N \rceil$  inter-patch graph layers, and fine-315 grained features of relatively densely sampled variables are extracted in the lower layers, while 316 coarse-grained features of both densely and sparsely sampled variables are extracted in the upper layers. The node embedding of the v-th variable is  $\mathcal{H}_v^{(\lceil \log_2 N \rceil)} \in \mathbb{R}^{d_{\text{model}}}$ . 317 318

- 319 320 4.5
  - 4.5 TASK DECODER

# 321 4.5.1 CLASSIFICATION

We simply calculate the sum of  $d_{\text{model}}$  channels for each variable's node embedding to get a V-dimensional vector C and use this to predict the final classification probabilities:  $\hat{y} =$  Softmax $(W^yC + b^y)$ . The training objective is to minimize the cross-entropy loss between  $\hat{y}$  and the ground truth y.

## 4.5.2 FORECASTING

Given a query  $q_j = (t_j, v_j)$ , we follow (Zhang et al., 2024) by concatenating the node embedding of the query variable  $v_j$  with the embedding of the query time  $t_j$ , and then pass it through an MLP projection layer to generate the forecasting result:  $\hat{z}_j = \text{MLP}([\mathcal{H}_{v_j}^{(\lceil \log_2 N \rceil)} \parallel \phi(t_j)])$ . The training objective is to minimize the mean squared error loss between  $\hat{z}_j$  and the ground truth  $z_j$ . The pseudo-code for Hi-Patch is presented in Appendix A (Algorithm 1).

## 5 EXPERIMENT

In this section, we present forecasting and classification experiments using a range of models and 8 datasets.

339 340

335

336 337

338

327

328

5.1 EXPERIMENTAL SETTING

341 342 343

## 5.1.1 DATASETS AND BASELINES

344 For the forecasting task, we follow (Zhang et al., 2024) and use four datasets: PhysioNet (Silva 345 et al., 2012), MIMIC-III (Johnson et al., 2016), Human Activity, and USHCN (Menne et al., 2015), 346 covering the fields of healthcare, biomechanics, and climate science. We compare our method with 347 seventeen relevant baselines, covering the SOTA models from (1) MTS forecasting: iTransformer 348 (Liu et al., 2024), ModernTCN(Luo & Wang, 2024), TimesNet (Wu et al., 2023), PatchTST (Nie 349 et al., 2023), (2) multi-scale MTS forecasting: Pathformer (Chen et al., 2024a), TimeMixer (Wang et al., 2024b), MSGNet (Cai et al., 2024), MICN (Wang et al., 2023), (3) IMTS classification: 350 Warpformer (Zhang et al., 2023), Raindrop (Zhang et al., 2022), GRU-D (Che et al., 2018), (4) IMTS 351 forecasting: tPatchGNN (Zhang et al., 2024), GraFITi (Yalavarthi et al., 2024), CRU (Schirmer et al., 352 2022), mTAND (Shukla & Marlin, 2021), Neural Flows (Biloš et al., 2021), Latent ODEs (Rubanova 353 et al., 2019). 354

355 For the classification task, we conduct experiments on four datasets in medical field where IMTS is most widely used, namely P19 (Reyna et al., 2020), PhysioNet (Silva et al., 2012), MIMIC-III 356 (Johnson et al., 2016) and P12 (Goldberger et al., 2000) where PhysioNet is a reduced version of 357 P12 considered by prior work (Shukla & Marlin, 2021). We compare our method with the state-of-358 the-art methods for irregular time series classification, including GRU-D (Che et al., 2018), ODE-359 RNN (Rubanova et al., 2019), IP-Net (Shukla & Marlin, 2019), SeFT (Horn et al., 2020), mTAND 360 (Shukla & Marlin, 2021), Raindrop (Zhang et al., 2022), StraTS (Tipirneni & Reddy, 2022), DuETT 361 (Labach et al., 2023), ViTST (Li et al., 2024) and Warpformer (Zhang et al., 2023). In addition, 362 we also compare our method with two approaches initially designed for forecasting tasks, namely 363 DGM<sup>2</sup>-O (Wu et al., 2021) and MTGNN (Wu et al., 2020). The implementation and hyperparameter 364 settings of these baselines are kept consistent with the optimal settings of the original paper. More 365 details of datasets and baselines can be found in Appendix B and C.

366 367

368

## 5.1.2 EVALUATION SETUP

For the forecasting task, we follow the data pre-processing method described in (Zhang et al., 2024) and randomly divide all the instances among each dataset into training, validation, and test sets according to ratios of 6:2:2. We use Mean Square Error (MSE) and Mean Absolute Error (MAE) to evaluate forecasting performance.

For the classification task, we follow the method described in (Harutyunyan et al., 2019) and divide the dataset into three parts for training, validation, and testing with the ratio of 70%,15%,15% on the MIMIC-III dataset. For the remaining three datasets, we adhered to (Zhang et al., 2022)'s approaches, and the ratio of training, validation, and testing set is 8:1:1. We measure the classification performance with the Area Under the Receiver Operating Characteristic Curve (AUROC) and Area Under the Precision-Recall Curve (AUPRC) since all the four datasets are binary classifi378 cation datasets with highly imbalanced class distribution. More details of metrics can be found in 379 Appendix D. 380

381 382

384

385

386

387

5.1.3 IMPLEMENTATION DETAILS

We adopt the Adam (Kingma & Ba, 2014) optimizer with a learning rate of 0.001, stopping it when the validation loss doesn't decrease over 10 epochs. All experiments are conducted with five random seeds, and the average and standard deviation are reported. All the models are experimented using the PyTorch library on 2 GeForce RTX-3090-24G GPUs. The detailed settings of hyperparameters can be found in Appendix E.

388 389 390

391

5.2 MAIN RESULTS

392 For the forecasting task, we test the model's performance under 3 varying observations and forecast 393 horizons on each dataset. Table 1 shows the results of the default horizon, and the complete results are presented in Table 4, 5, 6 and 7 (Appendix F.1). Table 2 reports the models' classification 394 performance on the other four datasets. In addition, to demonstrate the robustness of our method, 395 we test whether Hi-Patch can achieve good classification performance when a subset of variables is 396 completely missing. We uniformly discard 10%, 20%, 30%, 40%, and 50% of the variables, hiding 397 all their observations in both validation and test sets. Table 8 (Appendix F.2) reports the results. In 398 summary, Hi-Patch achieves the best performance in 62 out of 72 metrics across 8 datasets for both 399 classification and forecasting tasks. 400

Furthermore, regular MTS forecasting models, including single-scale and multi-scale approaches, 401 do not demonstrate competitive performance in IMTS forecasting. This is because regular MTS 402 methods have limited capability to handle the irregularities within and between variables in IMTS, 403 leading to suboptimal performance when directly applied to IMTS. Among methods specifically 404 designed for IMTS, Warpformer is the only method that considers the multi-scale characteristics 405 of IMTS and it achieves the most competitive performance. However, Warpformer involves in-406 terpolation of sparse variables, which may distort the original distribution of IMTS. In addition, 407 Warpformer employs attention separately in the time and variable dimensions, limiting its ability to 408 capture asynchronous dependencies. In contrast, our Hi-Patch flexibly represents and extracts both 409 synchronous and asynchronous dependencies of IMTS under different scales through an intra-patch 410 graph layer and several inter-patch graph layers. Moreover, Hi-Patch only handles actually observed points, avoiding the accumulation of imputation errors, particularly in cases of higher missing ratios, 411 exhibiting a degree of robustness. We also provide efficiency analysis in Appendix F.3. 412

413 414

417

Table 1: Method benchmarking on IMTS forecasting. The best results are highlighted in **bold**, and 415 the second-best results are in underlined. The results in the table are presented in the form of (Mean 416 ± Std). '-' indicates numerical overflow error.

18	Methods	Human Activity (	$3000 \text{ms} \rightarrow 1000 \text{ms})$	USHCN (24mor	$aths \rightarrow 1month$ )	PhysioNet (	$24h \rightarrow 24h$ )	MIMIC-III (	$24h \rightarrow 24h$ )
19		$MSE \times 10^{-3}$	$MAE \times 10^{-2}$	MSE×10 <sup>-1</sup>	$MAE \times 10^{-1}$	$MSE \times 10^{-3}$	$MAE \times 10^{-2}$	MSE×10 <sup>-2</sup>	$MAE \times 10^{-2}$
120 121	iTransformer ModernTCN TimesNet PatchTST	$\begin{array}{c} 3.97 \pm 0.10 \\ 3.99 \pm 0.05 \\ 3.79 \pm 0.05 \\ 5.21 \pm 0.33 \end{array}$	$\begin{array}{c} 4.30 \pm 0.08 \\ 4.32 \pm 0.04 \\ 4.28 \pm 0.04 \\ 5.10 \pm 0.20 \end{array}$		$\begin{array}{c} 4.18 \pm 0.55 \\ 3.58 \pm 0.09 \\ 3.56 \pm 0.12 \\ 3.66 \pm 0.13 \end{array}$	$53.55 \pm 19.59 \\ 28.99 \pm 11.06 \\ 9.30 \pm 0.70 \\ 25.56 \pm 4.42$	$\begin{array}{c} 16.87 \pm 4.51 \\ 5.94 \pm 0.20 \\ 5.50 \pm 0.34 \\ 10.90 \pm 1.08 \end{array}$	$7.24 \pm 0.50 \\10.07 \pm 9.79 \\2.34 \pm 0.54 \\7.24 \pm 0.65$	$21.24 \pm 0.97$ 10.76 ± 1.10 $8.09 \pm 0.10$ 19.80 ± 0.87
422 423 424	Pathformer TimeMixer MSGNet MICN	$3.40 \pm 0.16$ $4.97 \pm 0.31$ $6.32 \pm 0.16$ $6.93 \pm 0.12$	$\begin{array}{c} 3.65 \pm 0.08 \\ 5.02 \pm 0.16 \\ 6.06 \pm 0.09 \\ 6.04 \pm 0.06 \end{array}$	$ \begin{array}{c} - \\ 5.88 \pm 0.10 \\ 5.75 \pm 0.06 \\ 5.99 \pm 0.11 \end{array} $	- 3.59 ± 0.07 3.64 ± 0.09 3.69 ± 0.08	$\begin{array}{c} 6.75 \pm 0.41 \\ 13.98 \pm 0.31 \\ 9.84 \pm 0.29 \\ 10.34 \pm 0.24 \end{array}$	$\begin{array}{c} 4.61 \pm 0.20 \\ 6.88 \pm 0.09 \\ 5.79 \pm 0.10 \\ 6.00 \pm 0.14 \end{array}$	$- 4.78 \pm 0.09 \\ 2.65 \pm 0.18 \\ 2.36 \pm 0.06$	- 14.29 ± 0.06 9.10 ± 0.27 8.43 ± 0.11
425 426	Warpformer Raindrop GRU-D	$\frac{2.61 \pm 0.02}{4.42 \pm 0.25}$ $3.94 \pm 0.29$	$\frac{3.12 \pm 0.01}{4.65 \pm 0.14}$ $4.37 \pm 0.21$	$5.19 \pm 0.06 \\ 5.64 \pm 0.10 \\ 5.17 \pm 0.06$	$3.12 \pm 0.10$ $3.29 \pm 0.03$ $3.21 \pm 0.05$	$\frac{5.04 \pm 0.14}{10.63 \pm 0.29}$ 5.76 ± 0.34	$\frac{3.72 \pm 0.06}{6.02 \pm 0.19}$ $4.53 \pm 0.15$	$\frac{1.76 \pm 0.30}{2.31 \pm 0.07}$ $2.35 \pm 0.06$	$\frac{7.27 \pm 0.15}{8.61 \pm 0.12}$ 8.34 ± 0.22
427 428 429	tPatchGNN GraFITi CRU mTAND NeuralFlow Latent-ODE	$\begin{array}{c} 2.79 \pm 0.09 \\ 3.20 \pm 0.50 \\ 3.03 \pm 0.04 \\ 3.14 \pm 0.09 \\ 4.29 \pm 0.63 \\ 3.32 \pm 0.10 \end{array}$	$\begin{array}{c} 3.24 \pm 0.06 \\ 3.56 \pm 0.32 \\ 3.60 \pm 0.04 \\ 3.71 \pm 0.06 \\ 4.61 \pm 0.43 \\ 3.91 \pm 0.08 \end{array}$	$\begin{array}{c} \underline{5.00 \pm 0.03} \\ 5.07 \pm 0.03 \\ 5.15 \pm 0.50 \\ 5.03 \pm 0.05 \\ 5.41 \pm 0.05 \\ 5.16 \pm 0.04 \end{array}$	$3.07 \pm 0.05$ $2.97 \pm 0.04$ $3.18 \pm 0.03$ $3.00 \pm 0.06$ $3.35 \pm 0.06$ $3.21 \pm 0.07$	$5.06 \pm 0.10$ $5.32 \pm 0.27$ $6.43 \pm 0.62$ $6.18 \pm 0.31$ $7.68 \pm 0.37$ $6.85 \pm 0.28$	$\begin{array}{c} 3.75 \pm 0.07 \\ 4.00 \pm 0.14 \\ 4.51 \pm 0.16 \\ 4.44 \pm 0.19 \\ 4.84 \pm 0.19 \\ 4.77 \pm 0.17 \end{array}$	$\begin{array}{c} 1.97 \pm 0.05 \\ 1.76 \pm 0.04 \\ 2.23 \pm 0.03 \\ 2.15 \pm 0.05 \\ 2.34 \pm 0.05 \\ 2.11 \pm 0.15 \end{array}$	$\begin{array}{c} 7.76 \pm 0.22 \\ 7.34 \pm 0.23 \\ 7.99 \pm 0.22 \\ 8.00 \pm 0.06 \\ 8.09 \pm 0.09 \\ 7.76 \pm 0.08 \end{array}$
430	Hi-Patch	$2.57\pm0.02$	3.11 ± 0.03	4.94 ± 0.05	$2.96\pm0.04$	$4.86 \pm 0.03$	$3.62\pm0.07$	1.75 ± 0.26	$7.24 \pm 0.18$

Methods	P	19	Physi	ioNet	MIM	IC-III	P	12
	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC
GRU-D	88.7 ± 1.2	57.6 ± 2.3	79.1 ± 6.9	42.7 ± 7.2	82.2 ± 1.8	43.3 ± 2.1	79.6 ± 0.6	$41.7 \pm 1.8$
ODE-RNN	87.1 ± 1.0	$52.6 \pm 3.2$	$75.5 \pm 2.8$	$33.7 \pm 4.1$	$81.0 \pm 0.6$	$42.3 \pm 0.7$	$78.8 \pm 0.6$	$37.4 \pm 2.6$
IP-Net	$90.2 \pm 0.2$	$58.6 \pm 0.8$	$86.8 \pm 0.6$	$55.8 \pm 1.4$	$84.1 \pm 0.1$	$47.1 \pm 0.9$	$83.7 \pm 0.3$	$46.3 \pm 1.3$
SeFT	$84.0 \pm 0.3$	$49.3 \pm 0.5$	$75.5 \pm 0.2$	$29.4 \pm 0.9$	$67.9 \pm 0.2$	$\overline{23.2 \pm 0.4}$	$78.1 \pm 0.5$	$35.9 \pm 0.8$
MTGNN	$88.5 \pm 1.0$	$55.8 \pm 1.5$	77.1 ± 4.4	$35.4 \pm 7.3$	$78.5 \pm 2.3$	$35.2 \pm 3.1$	$82.1 \pm 1.5$	$41.8 \pm 2.1$
mTAND	$82.9 \pm 0.9$	$32.2 \pm 1.5$	$86.8 \pm 1.3$	$52.5 \pm 1.3$	$83.8 \pm 0.3$	$46.6 \pm 0.5$	$85.3 \pm 0.3$	$49.3 \pm 1.0$
DGM <sup>2</sup> -O	$91.6 \pm 0.5$	$60.0 \pm 1.3$	$85.8 \pm 0.7$	$50.4 \pm 3.2$	$80.4 \pm 0.5$	$36.0 \pm 0.8$	$85.8 \pm 0.1$	$48.3 \pm 0.7$
Raindrop	87.6 ± 2.7	$61.1 \pm 1.4$	$81.2 \pm 0.9$	$37.3 \pm 2.0$	79.8 ± 1.3	$35.2 \pm 1.1$	$82.0 \pm 0.6$	$42.7 \pm 1.7$
StraTS	$91.2 \pm 0.3$	$58.4 \pm 1.4$	84.9 ± 1.5	$47.3 \pm 5.3$	$84.4 \pm 0.4$	$46.4 \pm 0.8$	$86.7 \pm 0.7$	$52.1 \pm 1.5$
DuETT	$88.2 \pm 0.5$	$56.0 \pm 3.9$	$81.3 \pm 1.4$	$44.9 \pm 1.4$	$78.8 \pm 0.8$	$34.3 \pm 1.0$	$83.4 \pm 1.2$	$45.4 \pm 1.5$
ViTST	$91.7 \pm 0.1$	$57.5 \pm 0.7$	$81.3 \pm 1.9$	$37.4 \pm 2.9$	$81.8 \pm 0.3$	39.6 ± 1.3	$86.3 \pm 0.1$	$50.8 \pm 1.5$
Warpformer	$\underline{91.8\pm0.4}$	$60.6 \pm 2.6$	$83.3 \pm 0.7$	$43.5\pm2.3$	$84.6 \pm 0.3$	$46.6 \pm 0.9$	$85.4\pm0.5$	$50.4 \pm 1.5$
Hi-Patch	$92.1 \pm 0.4$	61.1 ± 2.1	86.8 ± 0.9	57.3 ± 1.9	84.8 ± 0.2	$47.2 \pm 1.0$	86.9 ± 0.7	53.3 ± 0.9

Table 2: Method benchmarking on IMTS classification. The best results are highlighted in **bold**, and the second-best results are in <u>underlined</u>. The results in the table are presented in the form of (Mean  $\pm$  Std %).

Table 3: Ablation results of Hi-Patch on four datasets. The results in the table are presented in the form of (Mean  $\pm$  Std).

Methods	Hunam Activity		USHCN		PhysioNet		MIMIC-III	
	MSE×10 <sup>-3</sup>	$MAE \times 10^{-2}$	MSE×10 <sup>-1</sup>	$MAE \times 10^{-1}$	$MSE \times 10^{-3}$	$MAE \times 10^{-2}$	$MSE \times 10^{-2}$	$MAE \times 10^{-2}$
Hi-Patch	2.57 ± 0.02	$3.11\pm0.03$	4.94 ± 0.05	$2.96 \pm 0.04$	$4.86 \pm 0.03$	$3.62\pm0.07$	$1.75\pm0.26$	$7.24\pm0.18$
w/o Hie w/o DVDT w/o 3W w/o TEAGG	$2.70 \pm 0.04 2.68 \pm 0.04 2.60 \pm 0.01 2.77 \pm 0.02$	$3.13 \pm 0.03$ $3.12 \pm 0.01$ $3.17 \pm 0.03$ $3.22 \pm 0.03$	$5.35 \pm 0.13 5.21 \pm 0.03 4.99 \pm 0.03 5.23 \pm 0.04$	$3.26 \pm 0.13$ $3.11 \pm 0.12$ $3.08 \pm 0.10$ $3.13 \pm 0.03$	$\begin{array}{c} 4.96 \pm 0.07 \\ 4.88 \pm 0.05 \\ 4.98 \pm 0.07 \\ 6.34 \pm 0.60 \end{array}$	$3.68 \pm 0.08$ $3.67 \pm 0.04$ $3.74 \pm 0.04$ $4.44 \pm 0.36$	$\begin{array}{c} 1.77 \pm 0.30 \\ 1.80 \pm 0.04 \\ 1.79 \pm 0.04 \\ 1.92 \pm 0.07 \end{array}$	$7.30 \pm 0.13$ $7.41 \pm 0.18$ $7.38 \pm 0.09$ $7.63 \pm 0.26$

5.3 ABLATION STUDY

In this section, we investigate the performance benefits generated by each key component of the proposed method on forecasting task. We compare the Hi-Patch with its four variants: (1) w/o Hie: We remove the hierarchical multi-scale architecture and set the patch size to the time span of the entire historical window, only extracting features at the original scale. (2) w/o DVDT: We removed the asynchronous edges of different variables at different times, retaining only the edges between nodes of the same variable at different times (SVDT) and different variables at the same time (DVST); (3) w/o 3W: We use only one set of attention parameter matrices  $\{W_q, W_k, W_v\}$ for three types of edges; (4) w/o TEAGG: We aggregate nodes with the same variables within a patch using mean aggregation rather than multi-time attention aggregation. The results demonstrate that all model components are necessary. w/o DVDT and w/o 3W demonstrate the importance of fully representing and capturing the temporal dependencies as well as both synchronous and asynchronous correlations among variables in IMTS. w/o TEAGG shows the effectiveness of multi-time attention in feature aggregation at the patch level, while w/o Hie highlights the necessity of considering the multi-scale characteristics in IMTS.

5.4 EFFECT OF SCALE QUANTITY





Figure 3: Effect of different scale quantities on four datasets.



Figure 4: Distribution of sample length and time span on four datasets. Each blue dot represents a sample, with its x-coordinate indicating the sample length, and the left y-axis representing its time span. The red curve is the distribution curve of sample length, corresponding to the right y-axis, representing the proportion of samples at a certain length.

Figure 3 illustrates the impact of different scale quantities across four datasets. Generally, increasing the number of scales improves performance within a certain range. However, beyond a certain point, further increasing the number of scales has a negative effect. This is because the number of scales is inversely related to patch size in our method: as the number of scales increases, the patch size decreases. A Patch that is too small may not contain enough observations to extract local patterns effectively. Additionally, we observe that on the PhysioNet and MIMIC-III datasets, performance with too many scales is even worse than with a single scale. To further investigate, we visualize the distribution of sample lengths and time spans across the four datasets in Figure 4. We find that samples in PhysioNet and MIMIC-III exhibit a 'short length, long span' characteristic, which indicates that samples in these two datasets predominantly show coarse-grained patterns with few fine-grained local features. In cases where the samples themselves contain few scale levels of features, using too many scales becomes redundant and can significantly degrade performance. We provide visualizations of the multi-scale views at different layers in our model in Appendix G.

510 511 512

513

494

495

496

497 498 499

500

501

502

503

504

505

506

507

508

509

#### 5.5 EFFECT OF PATCH SIZE

514 The impact of patch size on our model's per-515 formance on two real-world datasets, USHCN (climate) and PhysioNet (clinical), is illustrated 516 in Figure 5. For the USHCN dataset, the opti-517 mal model performance is achieved at the patch 518 size of 1.5 months. At this patch size, our hier-519 archical architecture extracts features at scales 520 of 1.5 months, 3 months, 6 months, and 12 521 months, which precisely cover several impor-522 tant observational scales in climatology, includ-523 ing the seasonal (Doblas-Reyes et al., 2013) 524 cycle (3 months), monsoonal (Clift & Plumb,



Figure 5: Effect of different patch size

525 2008) cvcle (6 months) and annual (Almazroui et al., 2012) cvcle (12 months). For the PhysioNet 526 dataset, the optimal model performance is achieved at the patch size of 6 hours. At this patch size, our model extracts features at 6-hour, 12-hour and 24-hour scales, which align with actual cycles 527 in clinical medicine. Among them, 6 hours is a common clinical monitoring period used in medi-528 cal practice (Seymour et al., 2017), while 12-hour and 24-hour cycles reflect circadian rhythms and 529 daily cycles, which are crucial for assessing patients' physiological changes and disease fluctuations 530 (Klerman et al., 2022). This alignment with real-world periodicity enables it to better capture the 531 inherent temporal dynamics in diverse applications. By aligning with critical observational and di-532 agnostic timeframes, our model enhances its predictive power and interpretability, making it highly 533 adaptable and effective across practical scenarios that require nuanced temporal understanding. 534

535

#### 6 CONCLUSION

536 537

538 In this paper, we propose Hi-Patch for modeling IMTS. The proposed method leverages the intra-539 patch/inter-patch graph neural network to flexibly represent and fully extract features at specific scales in IMTS. Based on this, the hierarchical architecture is used to effectively achieve multi-scale modeling of IMTS in a bottom-up manner (from local features to global ones). Our experimental results demonstrate that Hi-Patch outperforms existing methods in both IMTS forecasting and classification tasks. Our future work will focus on adaptive multi-scale modeling of IMTS, which selects the most suitable scales based on the specific temporal characteristics and dynamics of each sample to further improve performance.

#### References

546

547

551

552

553

554

576

577

578

579

580 581

582

583

584

- Mansour Almazroui, M Nazrul Islam, H Athar, PD Jones, and M Ashfaqur Rahman. Recent climate
   change in the arabian peninsula: annual rainfall and temperature analysis of saudi arabia for
   1978–2009. *International Journal of Climatology*, 32(6):953–966, 2012.
  - Marin Biloš, Johanna Sommer, Syama Sundar Rangapuram, Tim Januschowski, and Stephan Günnemann. Neural flows: Efficient alternative to neural odes. *Advances in Neural Information Processing Systems*, 34:21325–21337, 2021.
- Wanlin Cai, Yuxuan Liang, Xianggen Liu, Jianshuai Feng, and Yuankai Wu. Msgnet: Learning
   multi-scale inter-series correlations for multivariate time series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 11141–11149, 2024.
- Cristian Challu, Kin G Olivares, Boris N Oreshkin, Federico Garza Ramirez, Max Mergenthaler Canseco, and Artur Dubrawski. Nhits: Neural hierarchical interpolation for time series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pp. 6989–6997, 2023.
- 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.
- Peng Chen, Yingying Zhang, Yunyao Cheng, Yang Shu, Yihang Wang, Qingsong Wen, Bin Yang, and Chenjuan Guo. Pathformer: Multi-scale transformers with adaptive pathways for time series forecasting. In *International Conference on Learning Representations*, 2024a.
- Yuqi Chen, Kan Ren, Yansen Wang, Yuchen Fang, Weiwei Sun, and Dongsheng Li. Contiformer:
   Continuous-time transformer for irregular time series modeling. *Advances in Neural Information Processing Systems*, 36, 2024b.
- Zipeng Chen, Qianli Ma, and Zhenxi Lin. Time-aware multi-scale rnns for time series modeling. In
   *IJCAI*, pp. 2285–2291, 2021.
  - Peter D Clift and R Alan Plumb. *The Asian monsoon: causes, history and effects*, volume 288. Cambridge University Press Cambridge, 2008.
  - Jesse Davis and Mark Goadrich. The relationship between precision-recall and roc curves. volume 06, 06 2006. doi: 10.1145/1143844.1143874.
  - Edward De Brouwer, Jaak Simm, Adam Arany, and Yves Moreau. Gru-ode-bayes: Continuous modeling of sporadically-observed time series. *Advances in Neural Information Processing Systems*, 32, 2019.
- Francisco J Doblas-Reyes, Javier García-Serrano, Fabian Lienert, Aida Pintó Biescas, and Luis RL
   Rodrigues. Seasonal climate predictability and forecasting: status and prospects. *Wiley Interdis- ciplinary Reviews: Climate Change*, 4(4):245–268, 2013.
- Ary L Goldberger, Luis AN Amaral, Leon Glass, Jeffrey M Hausdorff, Plamen Ch Ivanov, Roger G Mark, Joseph E Mietus, George B Moody, Chung-Kang Peng, and H Eugene Stanley. Physiobank, physiotoolkit, and physionet: components of a new research resource for complex physiologic signals. *circulation*, 101(23):e215–e220, 2000.
- 593 Hrayr Harutyunyan, Hrant Khachatrian, David C Kale, Greg Ver Steeg, and Aram Galstyan. Multitask learning and benchmarking with clinical time series data. *Scientific data*, 6(1):96, 2019.

622

632

633

634

635

- Max Horn, Michael Moor, Christian Bock, Bastian Rieck, and Karsten Borgwardt. Set functions for time series. In *International Conference on Machine Learning*, pp. 4353–4363. PMLR, 2020.
- Sheo Yon Jhin, Minju Jo, Taeyong Kong, Jinsung Jeon, and Noseong Park. Ace-node: Attentive
   co-evolving neural ordinary differential equations. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pp. 736–745, 2021.
- Ming Jin, Huan Yee Koh, Qingsong Wen, Daniele Zambon, Cesare Alippi, Geoffrey I. Webb, Irwin King, and Shirui Pan. A survey on graph neural networks for time series: Forecasting, classification, imputation, and anomaly detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12):10466–10485, 2024.
- Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad
   Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. Mimic-iii,
   a freely accessible critical care database. *Scientific data*, 3(1):1–9, 2016.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Elizabeth B Klerman, Allison Brager, Mary A Carskadon, Christopher M Depner, Russell Foster, Namni Goel, Mary Harrington, Paul M Holloway, Melissa P Knauert, Monique K LeBourgeois, et al. Keeping an eye on circadian time in clinical research and medicine. *Clinical and Translational Medicine*, 12(12):e1131, 2022.
- Alex Labach, Aslesha Pokhrel, Xiao Shi Huang, Saba Zuberi, Seung Eun Yi, Maksims Volkovs, Tomi Poutanen, and Rahul G Krishnan. Duett: dual event time transformer for electronic health records. In *Machine Learning for Healthcare Conference*, pp. 403–422. PMLR, 2023.
- Kiang Lan, Hanshu Yan, Shenda Hong, and Mengling Feng. Towards enhancing time series contrastive learning: A dynamic bad pair mining approach. In *International Conference on Learning Representations*, 2024.
- Zekun Li, Shiyang Li, and Xifeng Yan. Time series as images: Vision transformer for irregularly
   sampled time series. Advances in Neural Information Processing Systems, 36, 2024.
- Zhonghang Li, Lianghao Xia, Yong Xu, and Chao Huang. Flashst: A simple and universal prompttuning framework for traffic prediction. In *Forty-first International Conference on Machine Learning*.
- Shizhan Liu, Hang Yu, Cong Liao, Jianguo Li, Weiyao Lin, Alex X Liu, and Schahram Dustdar.
   Pyraformer: Low-complexity pyramidal attention for long-range time series modeling and forecasting. In *International Conference on Learning Representations*, 2022.
  - Yong Liu, Tengge Hu, Haoran Zhang, Haixu Wu, Shiyu Wang, Lintao Ma, and Mingsheng Long. itransformer: Inverted transformers are effective for time series forecasting. In *International Conference on Learning Representations*, 2024.
- Donghao Luo and Xue Wang. Moderntcn: A modern pure convolution structure for general time
   series analysis. In *International Conference on Learning Representations*, 2024.
- Liantao Ma, Junyi Gao, Yasha Wang, Chaohe Zhang, Jiangtao Wang, Wenjie Ruan, Wen Tang, Xin Gao, and Xinyu Ma. Adacare: Explainable clinical health status representation learning via scaleadaptive feature extraction and recalibration. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 825–832, 2020.
- MJ Menne, CN Williams Jr, RS Vose, and Data Files. Long-term daily climate records from stations across the contiguous united states, 2015.
- Yuqi Nie, Nam H. Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A time series is worth
   64 words: Long-term forecasting with transformers. In *International Conference on Learning Representations*, 2023.

686

687

688

689

- Saeed Rahmani, Asiye Baghbani, Nizar Bouguila, and Zachary Patterson. Graph neural networks for intelligent transportation systems: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 24(8):8846–8885, 2023.
- Matthew A Reyna, Christopher S Josef, Russell Jeter, Supreeth P Shashikumar, M Brandon West-over, Shamim Nemati, Gari D Clifford, and Ashish Sharma. Early prediction of sepsis from clinical data: the physionet/computing in cardiology challenge 2019. *Critical care medicine*, 48 (2):210–217, 2020.
- 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 *International Conference on Machine Learning*, pp. 19388–19405. PMLR, 2022.
- 662
  663
  664
  664
  665
  666
  666
  666
  667
  668
  668
  668
  669
  669
  660
  660
  660
  660
  660
  661
  662
  662
  662
  663
  664
  664
  665
  665
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
  666
- Mohammad Amin Shabani, Amir H. Abdi, Lili Meng, and Tristan Sylvain. Scaleformer: Itera tive multi-scale refining transformers for time series forecasting. In *International Conference on Learning Representations*, 2023.
- Satya Narayan Shukla and Benjamin Marlin. Interpolation-prediction networks for irregularly sampled time series. In *International Conference on Learning Representations*, 2019.
- Satya Narayan Shukla and Benjamin Marlin. Multi-time attention networks for irregularly sampled time series. In *International Conference on Learning Representations*, 2021.
- Ikaro Silva, George Moody, Daniel J Scott, Leo A Celi, and Roger G Mark. Predicting in-hospital
   mortality of icu patients: The physionet/computing in cardiology challenge 2012. In 2012 com *puting in cardiology*, pp. 245–248. IEEE, 2012.
- Qingxiong Tan, Mang Ye, Grace Lai-Hung Wong, and Pong Chi Yuen. Cooperative joint attentive network for patient outcome prediction on irregular multi-rate multivariate health data. In *IJCAI*, pp. 1586–1592, 2021.
- Sindhu Tipirneni and Chandan K Reddy. Self-supervised transformer for sparse and irregularly
   sampled multivariate clinical time-series. *ACM Transactions on Knowledge Discovery from Data* (*TKDD*), 16(6):1–17, 2022.
  - Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. Advances in Neural Information Processing Systems, 30, 2017.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua
   Bengio. Graph attention networks. In *International Conference on Learning Representations*, 2018.
- Yogesh Verma, Markus Heinonen, and Vikas Garg. ClimODE: Climate and weather forecasting with physics-informed neural ODEs. In *International Conference on Learning Representations*, 2024.
- Huiqiang Wang, Jian Peng, Feihu Huang, Jince Wang, Junhui Chen, and Yifei Xiao. Micn: Multi scale local and global context modeling for long-term series forecasting. In *International Confer- ence on Learning Representations*, 2023.
- Jianian Wang, Sheng Zhang, Yanghua Xiao, and Rui Song. A review on graph neural network methods in financial applications. *arXiv preprint arXiv:2111.15367*, 2021.

- Lijing Wang, Aniruddha Adiga, Jiangzhuo Chen, Adam Sadilek, Srinivasan Venkatramanan, and Madhav Marathe. Causalgnn: Causal-based graph neural networks for spatio-temporal epidemic forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pp. 12191–12199, 2022.
- Shiyu Wang, Haixu Wu, Xiaoming Shi, Tengge Hu, Huakun Luo, Lintao Ma, James Y Zhang, and JUN ZHOU. Timemixer: Decomposable multiscale mixing for time series forecasting. In *International Conference on Learning Representations*, 2024a.
- Shiyu Wang, Haixu Wu, Xiaoming Shi, Tengge Hu, Huakun Luo, Lintao Ma, James Y Zhang, and Jun Zhou. Timemixer: Decomposable multiscale mixing for time series forecasting. *arXiv* preprint arXiv:2405.14616, 2024b.
- Haixu Wu, Tengge Hu, Yong Liu, Hang Zhou, Jianmin Wang, and Mingsheng Long. Timesnet: Temporal 2d-variation modeling for general time series analysis. In *International Conference on Learning Representations*, 2023.
- Yinjun Wu, Jingchao Ni, Wei Cheng, Bo Zong, Dongjin Song, Zhengzhang Chen, Yanchi Liu, Xuchao Zhang, Haifeng Chen, and Susan B Davidson. Dynamic gaussian mixture based deep generative model for robust forecasting on sparse multivariate time series. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 651–659, 2021.
- Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, Xiaojun Chang, and Chengqi Zhang. Connecting the dots: Multivariate time series forecasting with graph neural networks. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 753–763, 2020.
- 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.
- Jiawen Zhang, Shun Zheng, Wei Cao, Jiang Bian, and Jia Li. Warpformer: A multi-scale modeling approach for irregular clinical time series. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 3273–3285, 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*, 2024.
- Xiang Zhang, Marko Zeman, Theodoros Tsiligkaridis, and Marinka Zitnik. Graph-guided network
   for irregularly sampled multivariate time series. In *International Conference on Learning Representations*, 2022.

741 742 743

740

- 744 745
- 746 747 748

# 756 A Algorithm of Hi-Patch

Alg	orithm 1 The pseudo-code of Hi-Patch for forecasting
Inp	<b>ut</b> : An IMTS sample $S_i$ with M observations, a split time $t_S$ historical window $\mathcal{X}_i :=$
$\{(t)\}$	$ z_j, v_j   j = 1,, M, t_j \leq t_S$ , forecasting query $Q_i := \{[(t_j, v_j)]  j = 1,, M, t_j > t_S\}$ , ch size P, total time span T
Ou	<b>tput</b> : Predicted value set $\hat{\mathcal{Z}}_i$
1:	▷ Observation Encoder
2:	for $j = 1, 2,,  \mathcal{X}_i $ do
3:	encode observation tuple $o_i = (t_i, z_i, v_i)$ as graph node embedding $h_i$ using Eq.(5), get $\mathcal{H}^{(0)}$
4:	end for
5:	▷ Intra-Patch Graph Layer
6:	for $p = 1, 2,, \left\lceil \frac{T}{P} \right\rceil$ do
7:	construct intra-patch graph in patch $p$ using Eq.(6) and (7)
8:	update states of intra-patch graph nodes in patch $p$ through GAT using Eq.(8) and Eq.(9)
9:	aggregate nodes of the same variable in patch $p$ through multi-time attention using Eq.(10),
	get $\mathcal{H}^{(0)}$
10:	end for
11:	▷ Inter-Patch Graph Layers
12:	for $l = 1, 2,, \lceil \log_2 \left\lceil \frac{T}{P} \right\rceil \rceil$ do
13:	construct <i>P</i> -scale inter-patch graph
14:	update states of <i>P</i> -scale inter-patch graph nodes through GAT using Eq.(11)
15:	aggregate nodes of every two adjacent nodes of the same variable in P-scale inter-patch graph
	through multi-time attention using Eq.(12), get $\mathcal{H}^{(l)}$
16:	end for
17:	▷ Task Decoder
18:	for $j=1,2,, \mathcal{Q}_i $ do
19:	$\hat{z}_{j} = \mathrm{MLP}([\mathcal{H}_{v_{j}}^{(\lceil \log_{2} \left\lceil \frac{T}{P} \right\rceil \rceil)} \parallel \phi(t_{j})])$
20:	$\mathcal{Z}_i \leftarrow \mathcal{Z}_i \cup \{\hat{z}_j\}$
21:	end for _
22:	return $\mathcal{Z}_i$
R	Data sets
ע	
B.1	Forecasting
For	the forecasting task, we use four datasets and follow (Zhang et al., 2024)'s data preprocessing
pro	gram. Here is the detailed information of these datasets. We use three historical/forecasting

horizons on each dataset.
PhysioNet (Silva et al., 2012) This dataset includes 12000 IMTS from different patients, each with 41 clinical signals collected irregularly during the first 48 hours of ICU admission. We use the first 24/36/12 hours as the observed data to predict the queried values in the subsequent 24/12/36

799

hours.

800

MIMIC-III (Johnson et al., 2016) MIMIC-III is a clinical database containing IMTS data from 23457 patients, each with 96 variables recorded during the first 48 hours of ICU admission. We use the first 24/36/12 hours as the observed data to predict the queried values in the subsequent 24/12/36 hours.

Human Activity This dataset consists of 12 irregularly measured 3D positional variables from sensors worn on the ankles, belts, and chests of five individuals performing various activities. To better align with realistic forecasting scenarios, the original time series is chunked into 5400 IMTS, each spanning 4000 milliseconds. The first 3000/2000/1000 milliseconds are used as observed data to predict the sensor positions for the next 1000/2000/3000 milliseconds.

810 **USHCN** (Menne et al., 2015) The USHCN dataset includes over 150 years of climate data from 811 multiple U.S. stations, covering 5 climate variables. Following established preprocessing methods, 812 we focus on data from 1114 stations between 1996 and 2000, resulting in 26736 IMTS. Each in-813 stance uses data from the previous 24 months to predict the next 1/6/12 month's climate conditions. 814

815 **B**.2 CLASSIFICATION 816

817 P19 (Reyna et al., 2020) The PhysioNet Sepsis Early Prediction Challenge 2019 dataset contains 818 medical records of 38,803 patients. Each record includes 34 variables and a static vector detailing 819 attributes such as age, gender, time between hospital and ICU admission, ICU type, and ICU length 820 of stay in days. Each patient also receives a binary label indicating whether sepsis occurs within the next 6 hours. We exclude samples with excessively short or long time series following (Zhang et al., 821 2022). Available at https://physionet.org/content/challenge-2019/1.0.0/. 822

823

824 P12 (Goldberger et al., 2000) The P12 dataset comprises data from 11,988 patients after removing 825 12 inappropriate samples identified by (Horn et al., 2020). Each record includes multivariate time series data from the first 48 hours of ICU stay, consisting of 36 sensor measurements (excluding 826 weight) and a static vector with 9 elements, including age and gender. Patients are labeled based on 827 ICU stay duration: a negative label indicates three days or less, and a positive label indicates more 828 than three days. Available at https://physionet.org/content/challenge-2012/1.0.0/. 829

830

**MIMIC-III** (Johnson et al., 2016) MIMIC-III is a widely used dataset containing de-identified 831 EHRs of ICU patients admitted to Beth Israel Deaconess Medical Center from 2001 to 2012, orig-832 inally with around 57,000 records covering variables such as medications and vital signs. We fo-833 cus on the in-hospital mortality prediction task, using a subset established by (Harutyunyan et al., 834 2019). After preprocessing, our dataset includes 16 features and 21,107 data points. Available at 835 https://physionet.org/content/mimiciii/1.4/. 836

837 **PhysioNet** (Silva et al., 2012) Physionet contains the data from the first 48 hours of ICU patients, 838 which is a reduced version of P12 considered by prior work. Therefore, we follow the same prepro-839 cessing methods used for the P12 dataset. The processed data set includes 3997 labeled instances. 840 We focus on predicting in-hospital mortality.

- 842 843 С
- 844 845

846 847

848

849

850

851

852

853 854

855

856 857

861

841

C.1 FORECASTING

BASELINES

C.1.1 METHODS FOR REGULAR MTS

For methods for regular MTS, we organize each sample into a  $V \times T_h$  history matrix and a  $V \times T_f$ forecasting matrix, where V represents the maximum number of variables and  $T_h/T_f$  represents the historical/forecasting length of the sample. Unobserved positions are filled with zeros, the sequence length is set to the maximum historical length across all samples, and the forecasting length is set to the maximum forecasting length across all samples.

iTransformer (Liu et al., 2024) uses inverted Transformers for time series forecasting. We use the following setting in our experiment:  $e_{\text{layers}} = 3$ ,  $d_{\text{layers}} = 1$ , factor = 3,  $d_{\text{model}} = 512$ ,  $d_{\text{ff}} = 512$ .

**TimesNet** (Wu et al., 2023) analyses general time series through temporal 2D-Variation. We use 858 the following setting in our experiment:  $e_{\text{lavers}} = 2$ ,  $d_{\text{lavers}} = 1$ , factor = 3,  $d_{\text{model}} = 256$ ,  $d_{\text{ff}} = 512$ , 859 top-k = 5.860

**PatchTST** (Nie et al., 2023) is a Transformer-based model using patch and channel independence 862 for long-term time series forecasting. We use the following setting in our experiment:  $e_{\text{layers}} = 2$ , 863  $d_{\text{lavers}} = 1$ , factor = 3, patch\_len = 16, stride = 8.

**MICN** (Wang et al., 2023) achieves multi-scale local and global context modeling for long-term time series forecasting. We use the following setting in our experiment:  $e_{\text{layers}} = 2$ ,  $d_{\text{layers}} = 1$ , factor = 3,  $d_{\text{model}} = 256$ ,  $d_{\text{ff}} = 512$ , top\_k = 5, decomp\_kernel = 32, conv\_kernel = 24, isometric\_kernel = [18, 6].

**TimeMixer** (Wang et al., 2024b) achieves complementary predictive capabilities by disentangling variations in multi-scale series. We use the following setting in our experiment:  $e_{\text{layers}} = 3$ ,  $d_{\text{layers}} = 1$ , factor = 3,  $d_{\text{model}} = 16$ ,  $d_{\text{ff}} = 32$ , down\_sampling\_layers = 3, down\_sampling\_window = 2.

For the above 5 methods, we use the implementation in Time-Series-Library

ModernTCN (Luo & Wang, 2024) is a modern pure convolution structure designed for general time series analysis. We use the implementation in https://github.com/luodhhh/ModernTCN and the following setting in our experiment: ffn\_ratio = 8, patch\_size = 8, patch\_stride = 4, num\_blocks = 1, large\_size = 51, small\_size = 5, dims = 64

878 879 880

881

882

883 884

885

886

887

888

889 890

891

895

901

906

907

908

909

869

870

871

**Pathformer** (Chen et al., 2024a) is a multi-scale transformer architecture with adaptive pathways for time series forecasting. We use the implementation in https://github.com/decisionintelligence/pathformer and the following setting in our experiment: k = 2, layer\_nums = 3,  $d_{\text{model}} = 16$ ,  $d_{\text{ff}} = 64$ .

**MSGNet** (Cai et al., 2024) learns multi-scale inter-series correlations for multivariate time series forecasting. We use the implementation in https://github.com/YoZhibo/MSGNet and the following setting in our experiment:  $e_{\text{layers}} = 2$ ,  $d_{\text{model}} = 512$ ,  $d_{\text{ff}} = 64$ , n\_heads = 8, top\_k = 5, dropout = 0.1, nums\_kernels = 6, conv\_channel = 32, skip\_channel = 32, gcn\_depth = 2, propalpha = 0.3, node\_dim = 10, gcn\_dropout = 10.

C.1.2 METHODS FOR IMTS

IMTS forecasting methods can directly make predictions. For IMTS classification methods, we use
 them as encoders to extract variable-level representations for each sample, followed by forecasting
 using the decoder described in Section 4.5.2.

896 Warpformer (Zhang et al., 2023) A transformer-based network that captures features at different 897 scales in IMTS using warping modules and dual attention mechanisms. We use three scales with 898 normalized length  $\tilde{L}^{(0)} = 0$ ,  $\tilde{L}^{(1)} = 0.2$  and  $\tilde{L}^{(2)} = 1$ . The dimension of representations *D* is set as 899 32. The attention heads and the layers of the warpformer are set as 1 and 2, respectively. We use the 900 implementation at https://github.com/imJiawen/Warpformer.

**Raindrop** (Zhang et al., 2022) A graph neural network that embeds IMTS while learning the dynamics of sensors purely from observation data. We use the following setting in our experiment:  $d_{ob} = 4$ ,  $p^t = 16$ ,  $r_v = 16$ , L = 2,  $d_k = 20$ ,  $d_a = V$ . We use the implementation at https://github.com/mims-harvard/Raindrop.

**GRU-D** (Che et al., 2018) GRU-D takes two representations of missing patterns, i.e., masking and time interval, and effectively incorporates them into a deep model architecture. The number of hidden states of GRU-D is set as 49. We use the implementation from https://github.com/Han-JD/GRU-D.

910 911

**tPatchGNN** (Zhang et al., 2024) is a transformable patching graph neural network for IMTS fore casting. We use the implementation in https://github.com/usail-hkust/t-PatchGNN and use the hyperparameters specified in their scripts.

**<sup>916</sup> GraFITi** (Yalavarthi et al., 2024) use bipartite graph for representing and forecasting of IMTS. **917** We use the following setting in our experiment: latent\_dim = 128, n\_layers = 4,  $n_{heads} = 1$ . We use the implementation at https://github.com/yalavarthivk/GraFITi.

CRU (Schirmer et al., 2022) models IMTS with continuous recurrent units. We use the following setting in our experiment: latent\_state\_dim = 20, hidden\_units = 50, bandwidth = 10, num\_basis = 20, trans\_covar = 0.1. We use the implementation at https://github.com/boschresearch/Continuous-Recurrent-Units.

mTAND (Shukla & Marlin, 2021) A deep learning framework for IMTS data that learns an embedding of continuous time values and uses an attention mechanism to produce a fixed-length representation. We set the latent dimension and the hidden size of GRU to 32. The number of reference points
and the dimension of time embedding is 128. We use the implementation at https://github.com/remllab/mTAN.

NeuralFlows (Biloš et al., 2021) use neural networks to model ODE solution curves to mitigate the expensive numerical solvers in neural ODEs. We use the implementation at https://github.com/mbilos/neural-flows-experiments and the following setting in our experiment: flow\_model='CouplingFlow', hidden\_dim = 50, hidden\_layers = 3, latents = 20, rec\_dims = 40.

Latent-ODEs (Rubanova et al., 2019) is an ODE-based model that improves RNNs with continuous-time hidden state dynamics specified by neural ODEs. We use the implementation in https://github.com/YuliaRubanova/latent\_ode and the following setting in our experiment: latents = 20, units = 50, gen\_layers = 3, rec\_dims = 40, rec\_layers = 3, gru\_units = 50.

938 C.2 CLASSIFICATION

937

939

954

969

ODE-RNN (Rubanova et al., 2019) ODE-RNN uses neural ODEs to model hidden state dynamics and an RNN to update the hidden state in the presence of a new observation. The latent dimension is set as 40, and the ODE function has 3 layers with 50 units. We use the implementation at https://github.com/YuliaRubanova/latent\_ode

**SEFT** (Horn et al., 2020) A set function approach where all the observations are modeled individually before pooling them together using an attention-based approach. We use a constant architecture for the attention network f' with 2 layers, 4 heads and dimensionality of the dot product space dof 128. In addition, the attention network f' was always set to use mean aggregation. We use the implementation from https://github.com/BorgwardtLab/SeFT.

**IP-Net** (Shukla & Marlin, 2019) A model architecture for IMTS data based on several semiparametric interpolation layers organized into an interpolation network followed by a prediction
network GRU. The number of reference points is set as 192. The hidden size of GRU is 100. We
take the source code at https://github.com/mlds-lab/interp-net.

DGM<sup>2</sup>-O (Wu et al., 2021) A generative model, which tracks the transition of latent clusters instead of isolated feature representations, achieves robust sparse time series modeling. We use the DGM<sup>2</sup>-O and set both the hidden dimension and the cluster\_num as 10. We use the source code at https://github.com/thuwuyinjun/DGM2.

MTGNN (Wu et al., 2020) A general graph neural network framework designed for MTS data.
We use 5 graph convolution and 5 temporal convolution modules with the dilation exponential factor 2. The graph convolution and temporal convolution modules have 16 output channels. The skip connection layers all have 32 output channels. The first layer of the output module has 64 output channels, and the second layer has 1 output channel. We use the implementation from: https://github.com/nnzhan/MTGNN.

StraTS (Tipirneni & Reddy, 2022) is a self-supervised transformer for sparse IMTS. We use the implementation at https://github.com/sindhura97/STraTS and the following setting in our experiment: hidden\_dim = 64, num\_layers = 2, num\_heads = 16, dropout = 0.2.

**DuETT** (Labach et al., 2023) is a dual event time transformer for Electronic Health Records (EHRs). We use the implementation at https://github.com/layer6ai-labs/DuETT and the default settings of the model declaration in this repository.

972 ViTST (Li et al., 2024): transforms IMTS into line graph images and adapts powerful vision transformers to perform time series classification in the same way as image classification. We use the implementation at https://github.com/Leezekun/ViTST.
975

Please refer to C.1.2 for details of GRU-D (Che et al., 2018), mTAND (Shukla & Marlin, 2021),
Raindrop (Zhang et al., 2022) and Warpformer (Zhang et al., 2023).

## D PERFORMANCE METRICS

**MSE** MSE (Mean Squared Error) measures the average of the squared differences between predicted and actual values. The calculation formula is:

$$MSE = \frac{1}{n} \sum_{i=1}^{n} (y_i - \hat{y}_i)^2,$$
(14)

where  $y_i$  represents the actual value,  $\hat{y}_i$  represents the predicted value, and n is the number of observations. A smaller MSE indicates better model performance. Since errors are squared, MSE is sensitive to large errors or outliers.

**MAE** MAE (Mean Absolute Error) measures the average of the absolute differences between predicted and actual values. The calculation formula is:

$$MAE = \frac{1}{n} \sum_{i=1}^{n} |y_i - \hat{y}_i|,$$
(15)

where  $y_i$  represents the actual value,  $\hat{y}_i$  represents the predicted value, and n is the number of observations. A smaller MAE indicates better model performance. Compared to MSE, MAE is less sensitive to outliers and provides a straightforward average measure of error.

AUROC AUROC is commonly employed in binary classification tasks, where one class is designated as positive and the other as negative. It represents the area under the Receiver Operating Characteristic (ROC) curve, constructed by plotting the True Positive Rate (TPR) against the False Positive Rate (FPR). AUROC ranges from 0 to 1, with a higher value indicating better model performance in accurately discriminating between positive and negative instances. An AUROC equal to 0.5 indicates a model's performance equivalent to random guessing, while an AUROC greater than 0.5 signifies superiority over random guessing.

AUPRC The Area Under the Precision-Recall Curve is widely used as a performance metric for imbalanced binary classification tasks. It provides a comprehensive assessment of a model's precision-recall trade-off. The Precision-Recall curve is constructed by plotting recall on the x-axis and precision on the y-axis. AUPRC ranges from 0 to 1, and a higher value indicates better model performance in achieving high precision and recall simultaneously. It has been suggested as a good criterion for unevenly distributed classification problems (Davis & Goadrich, 2006).

1012 1013

976

977 978

979 980

981

987

988

989 990

991

992 993

994 995

996

997

998

#### 1014 E HYPERPARAMETERS SETTINGS

1015 We search all hyperparameters in the grid to find the best hyperparameters for our proposed model 1016 Hi-Patch. Specifically, our model has a total of 4 hyperparameters: patch size P, dimension of 1017 node state  $d_{\text{model}}$ , number of multi-head attention heads  $n_{\text{heads}}$ , number of GAT layers L. Since 1018 the number of patches  $N = \left| \frac{T}{P} \right|$  and the number of inter-patch graph layers equals to  $\left[ \log_2 N \right]$ 1019 where T is the dataset-specific maximum time span, we search N over the range  $\{2, 4, 8, 16, 32\}$  to 1020 maintain the number of inter-patch graph layers as an integer. Thus the patch size P of each dataset 1021 is determined by  $P = \left| \frac{T}{N} \right|$ . Additionally, we search  $d_{\text{model}}$  in  $\{16, 32, 64, 128\}$ ,  $n_{\text{heads}}$  in  $\{1, 2, 4, 8\}$ 1022 and L in  $\{1, 2, 3\}$ . The best hyperparameters for each dataset are reported in the code.

1023

1024

# 1026 F ADDITIONAL EXPERIMENTS

# 1028 F.1 VARYING OBSERVATION AND FORECAST HORIZONS

Table 4: Performance of varying observation and forecast horizons on Human Activity dataset. The best results are highlighted in **bold**, and the second-best results are in <u>underlined</u>. The results in the table are presented in the form of (Mean ± Std).

Horizon	2000ms -	$\rightarrow 2000 \mathrm{ms}$	$1000 \text{ms} \rightarrow 3000 \text{ms}$			
Metric	$MSE \times 10^{-3}$	$MAE \times 10^{-2}$	$MSE \times 10^{-3}$	$MAE \times 10^{-2}$		
iTransformer	7.49 ± 4.72	$6.08 \pm 2.17$	$5.58 \pm 0.04$	$5.13 \pm 0.05$		
ModernTCN	$5.26 \pm 0.06$	$4.96 \pm 0.04$	$12.12 \pm 1.01$	$5.72 \pm 0.07$		
TimesNet	$5.38 \pm 0.30$	$5.32 \pm 0.20$	$9.90 \pm 0.42$	$7.34 \pm 1.81$		
PatchTST	$7.25 \pm 0.29$	$6.26 \pm 0.17$	$8.97 \pm 1.96$	$6.94\pm0.94$		
Pathformer	$4.67 \pm 0.22$	$4.58 \pm 0.17$	$5.49 \pm 0.12$	$5.05 \pm 0.06$		
TimeMixer	$5.39 \pm 0.54$	$5.05 \pm 0.38$	$5.96 \pm 0.19$	$5.36 \pm 0.10$		
MSGNet	$7.90 \pm 0.36$	$6.85 \pm 0.13$	$8.29 \pm 0.22$	$7.02 \pm 0.12$		
MICN	$7.57 \pm 0.05$	$6.43 \pm 0.02$	$8.16 \pm 0.12$	$6.78\pm0.05$		
Warpformer	$3.60 \pm 0.08$	$3.81 \pm 0.03$	$4.26 \pm 0.11$	$4.26 \pm 0.04$		
Raindrop	$5.57 \pm 0.34$	$5.15 \pm 0.11$	$5.75 \pm 0.33$	$5.37 \pm 0.22$		
GRU-D	$5.93 \pm 0.10$	$5.66 \pm 0.66$	$6.14 \pm 0.76$	$5.75 \pm 0.49$		
tPatchGNN	3.71 ± 0.20	$3.89 \pm 0.16$	$4.56 \pm 0.08$	$4.32 \pm 0.06$		
GraFITi	$4.59 \pm 0.04$	$4.45 \pm 0.04$	$4.91 \pm 0.07$	$4.62 \pm 0.03$		
CRU	$4.12 \pm 0.08$	$4.43 \pm 0.06$	$4.85 \pm 0.09$	$4.86 \pm 0.07$		
mTAND	$4.38 \pm 0.37$	$4.59 \pm 0.29$	$5.29 \pm 0.32$	$5.12 \pm 0.23$		
NeuralFlow	$5.47 \pm 0.49$	$5.35 \pm 0.28$	6.01 ± 0.91	$5.66 \pm 0.60$		
Latent-ODE	$5.04 \pm 0.46$	$5.11 \pm 0.29$	$5.48 \pm 0.21$	$5.33 \pm 0.14$		
Hi-Patch	$3.29 \pm 0.04$	$3.70\pm0.04$	$4.21 \pm 0.08$	$4.25 \pm 0.07$		

Table 5: Performance of varying observation and forecast horizons on USHCN dataset. The best results are highlighted in **bold**, and the second-best results are in <u>underlined</u>. The results in the table are presented in the form of (Mean  $\pm$  Std).

Horizon	24months	$\rightarrow$ 6months	$24 \text{months} \rightarrow 12 \text{months}$			
Metric	$MSE \times 10^{-1}$	$MAE \times 10^{-1}$	$MSE \times 10^{-1}$	MAE×10 <sup>-</sup>		
iTransformer	$6.05 \pm 0.03$	$3.89 \pm 0.14$	6.19 ± 0.12	$4.01 \pm 0.12$		
ModernTCN	$6.03 \pm 0.10$	$3.68 \pm 0.06$	$6.19 \pm 0.07$	$3.75 \pm 0.03$		
TimesNet	$5.68 \pm 0.05$	$3.66 \pm 0.06$	$5.84 \pm 0.06$	$3.78 \pm 0.10$		
PatchTST	$6.12 \pm 0.03$	$4.01\pm0.08$	$6.55 \pm 0.12$	$4.20 \pm 0.06$		
TimeMixer	$5.89 \pm 0.07$	$3.65 \pm 0.05$	6.06 ± 0.27	$3.76 \pm 0.04$		
MSGNet	$5.76 \pm 0.02$	$3.74 \pm 0.08$	$5.93 \pm 0.05$	$3.77 \pm 0.03$		
MICN	$6.02 \pm 0.03$	$3.85\pm0.04$	$6.00 \pm 0.03$	$3.85 \pm 0.03$		
Warpformer	$5.12 \pm 0.03$	$3.13 \pm 0.08$	5.10 ± 0.07	$3.13 \pm 0.12$		
Raindrop	$7.01 \pm 0.49$	$4.24 \pm 0.33$	$7.61 \pm 0.02$	$4.61 \pm 0.05$		
GRU-D	$5.29 \pm 0.09$	$3.34 \pm 0.09$	$5.36 \pm 0.12$	$3.25 \pm 0.07$		
tPatchGNN	$5.23 \pm 0.02$	$3.24 \pm 0.19$	$6.23 \pm 0.10$	$3.83 \pm 0.60$		
GraFITi	$5.12 \pm 0.14$	$3.09 \pm 0.10$	$5.01 \pm 0.03$	$3.14 \pm 0.06$		
CRU	$6.77 \pm 1.04$	$4.11 \pm 0.61$	$6.64 \pm 0.95$	$4.08 \pm 0.51$		
mTAND	$5.16 \pm 0.10$	$3.10 \pm 0.07$	$5.07 \pm 0.03$	$3.09 \pm 0.03$		
NeuralFlow	$5.52 \pm 0.05$	$3.46 \pm 0.05$	$5.48 \pm 0.37$	$3.56 \pm 0.37$		
Latent-ODE	$5.18 \pm 0.04$	$3.36 \pm 0.04$	$5.23 \pm 0.04$	$3.35 \pm 0.02$		
Hi-Patch	$5.07 \pm 0.17$	$3.06 \pm 0.06$	$5.02 \pm 0.05$	$3.07 \pm 0.08$		

1082 1083 1084	Table 6: Performance of varying observation and forecast horizons on PhysioNet dataset. The best results are highlighted in <b>bold</b> , and the second-best results are in <u>underlined</u> . The results in the table are presented in the form of (Mean $\pm$ Std).
1085	

Horizon	36h -	$\rightarrow$ 12h	$12h \rightarrow 36h$			
Metric	$MSE \times 10^{-3}$	$MAE \times 10^{-2}$	$MSE \times 10^{-3}$	$MAE \times 10^{-2}$		
iTransformer ModernTCN TimesNet PatchTST	$56.83 \pm 21.17 24.87 \pm 9.44 9.43 \pm 0.53 26.13 \pm 1.75$	$17.05 \pm 5.21 \\ 6.58 \pm 0.58 \\ 5.55 \pm 0.20 \\ 11.25 \pm 0.29$	$54.17 \pm 17.91 31.48 \pm 3.96 9.26 \pm 0.18 25.63 \pm 1.51$	$17.03 \pm 4.18 \\ 6.73 \pm 0.53 \\ 5.58 \pm 0.10 \\ 10.70 \pm 0.38$		
Pathformer TimeMixer MSGNet MICN	$\begin{array}{c} 6.85 \pm 0.42 \\ 12.52 \pm 0.45 \\ 10.44 \pm 0.44 \\ 10.98 \pm 0.31 \end{array}$	$\begin{array}{c} 4.61 \pm 0.18 \\ 6.47 \pm 0.21 \\ 5.94 \pm 0.08 \\ 6.01 \pm 0.11 \end{array}$	$\begin{array}{c} 8.19 \pm 0.28 \\ 19.86 \pm 0.17 \\ 9.91 \pm 0.12 \\ 10.24 \pm 0.14 \end{array}$	$5.03 \pm 0.07$ $8.38 \pm 0.18$ $5.74 \pm 0.08$ $5.85 \pm 0.06$		
Warpformer Raindrop GRU-D	$\frac{4.17 \pm 0.13}{10.67 \pm 0.33}$ 6.85 ± 0.37	$\frac{3.38 \pm 0.08}{5.87 \pm 0.20}$ $4.88 \pm 0.18$	$  \begin{array}{c} 6.51 \pm 0.12 \\ 10.24 \pm 0.18 \\ 7.80 \pm 0.22 \end{array} $	$4.24 \pm 0.04$ $5.83 \pm 0.10$ $5.13 \pm 0.13$		
tPatchGNN GraFITi CRU mTAND NeuralFlow Latent-ODE	$\begin{array}{c} 4.22 \pm 0.09 \\ 4.58 \pm 0.11 \\ 6.74 \pm 0.21 \\ 5.61 \pm 0.31 \\ 8.87 \pm 1.00 \\ 6.99 \pm 0.24 \end{array}$	$3.38 \pm 0.04  3.65 \pm 0.05  4.82 \pm 0.11  4.15 \pm 0.09  5.43 \pm 0.18  4.74 \pm 0.11$	$\begin{array}{c} 6.45 \pm 0.11 \\ \underline{6.30 \pm 0.14} \\ \overline{7.66 \pm 0.14} \\ \overline{7.46 \pm 0.19} \\ \overline{7.98 \pm 0.57} \\ \overline{7.28 \pm 0.13} \end{array}$	$\frac{4.24 \pm 0.09}{4.38 \pm 0.12}$ 4.97 ± 0.05 4.85 ± 0.05 5.08 ± 0.24 4.83 ± 0.07		
Hi-Patch	$4.16 \pm 0.08$	$3.31 \pm 0.06$	6.30 ± 0.06	$4.12 \pm 0.05$		

Table 7: Performance of varying observation and forecast horizons on MIMIC-III dataset. The best results are highlighted in **bold**, and the second-best results are in <u>underlined</u>. The results in the table are presented in the form of (Mean  $\pm$  Std). 

1112						
1113	Horizon	36h -	$\rightarrow$ 12h	$12h \rightarrow 36h$		
1114	Matria	$MSE \times 10^{-2}$	$MAE \times 10^{-2}$	$MSE \times 10^{-2}$	$MAE \times 10^{-2}$	
1115	Wietric	MSEX10	MAEXIU	MSEX10	MAEXIU	
1116	iTransformer	$7.27 \pm 0.04$	$21.52\pm0.09$	$7.44 \pm 0.21$	$21.47 \pm 0.45$	
1117	ModernTCN	$5.12 \pm 0.59$	$11.41 \pm 0.30$	$3.96 \pm 0.23$	$10.79 \pm 0.39$	
1118	TimesNet	$2.02 \pm 0.09$	$8.18 \pm 0.22$	$2.39 \pm 0.06$	$8.35 \pm 0.20$	
1110	Patch1S1	$8.81 \pm 0.96$	$22.97 \pm 1.32$	$/.11 \pm 0.65$	$19.96 \pm 1.24$	
1100	Pathformer	-	_	$2.45 \pm 0.02$	$8.65 \pm 0.09$	
1120	TimeMixer	$4.31 \pm 0.29$	$13.44 \pm 0.28$	$6.31 \pm 0.04$	$18.41 \pm 0.06$	
1121	MSGNet	$2.48 \pm 0.12$	$9.42 \pm 0.32$	$2.57 \pm 0.03$	$9.02 \pm 0.06$	
1122	MICN	$2.15 \pm 0.13$	$8.66 \pm 0.26$	$2.39 \pm 0.06$	$8.58 \pm 0.11$	
1123	Warpformer	$1.45 \pm 0.10$	$6.74 \pm 0.08$	$2.32 \pm 0.04$	$8.14 \pm 0.07$	
1124	Raindrop	$\overline{2.21 \pm 0.37}$	$\overline{9.17 \pm 0.49}$	$2.36 \pm 0.03$	$8.63 \pm 0.11$	
1125	GRU-D	$2.03 \pm 0.13$	$8.14\pm0.26$	$2.39\pm0.02$	$8.43 \pm 0.13$	
1126	tPatchGNN	1.44 + 0.08	$6.78 \pm 0.14$	$2.35 \pm 0.03$	8 23 + 0 08	
1127	GraFITi	$1.61 \pm 0.27$	$7.16 \pm 0.36$	$2.22 \pm 0.05$	$8.13 \pm 0.13$	
1128	CRU	$2.00 \pm 0.13$	$8.16 \pm 0.26$	$2.34 \pm 0.05$	$\overline{8.32 \pm 0.13}$	
1129	mTAND	$2.01 \pm 0.09$	$8.13 \pm 0.23$	$2.29 \pm 0.03$	$8.38 \pm 0.08$	
1130	NeuralFlow	$1.97 \pm 0.12$	$8.39 \pm 0.25$	$2.26 \pm 0.08$	$8.29 \pm 0.10$	
1100	Latent-ODE	$1.90 \pm 0.03$	$7.92 \pm 0.17$	$2.38 \pm 0.05$	$8.35 \pm 0.13$	
1101	Hi-Patch	$1.56 \pm 0.10$	6.71 ± 0.16	$2.32 \pm 0.02$	8.11 ± 0.08	
1132				1		
1133						

# 1134 F.2 LEAVE-VARIABLES-OUT CLASSIFICATION

Table 8: Classification performance on samples with a fixed set of left-out variables. The best results are highlighted in **bold** and the second best results are in <u>underlined</u>.

1143				Discard ratio									
1144	Dataset	Methods	10	)%	20	)%	30	1%	40	1%	50	)%	
1145			AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	
1146		GRU-D	$68.6 \pm 2.3$	$35.8 \pm 2.2$	68.2 ± 2.1	$34.5 \pm 2.9$	66.8 ± 3.3	$32.7 \pm 4.6$	$65.8 \pm 4.0$	$31.3 \pm 5.2$	$65.1 \pm 4.1$	$30.4 \pm 5.5$	
4 4 4 7		mTAND	$74.9 \pm 0.6$	$37.7 \pm 0.6$	$74.0 \pm 1.3$	$36.5 \pm 1.5$	$71.4 \pm 3.8$	$34.1 \pm 3.7$	$70.6 \pm 3.6$	$33.2 \pm 3.7$	$70.1 \pm 3.5$	$32.5 \pm 3.6$	
1147	P12	DGM <sup>-</sup> -O MTGNN	$76.3 \pm 1.1$	$39.3 \pm 1.5$ $20.5 \pm 1.5$	$70.1 \pm 1.1$	$38.2 \pm 1.7$	$74.8 \pm 2.2$	$36.8 \pm 2.6$ 28 5 ± 2.2	$72.0 \pm 5.3$	$34.3 \pm 5.0$	$70.4 \pm 5.9$	$32.7 \pm 5.7$	
1148	112	Raindron	$71.2 \pm 2.1$ 732 + 16	$30.3 \pm 1.3$ $32.4 \pm 0.9$	$70.3 \pm 3.3$ $73.0 \pm 1.6$	$29.7 \pm 2.0$ $31.7 \pm 1.1$	$722 \pm 726$	$28.5 \pm 3.3$ $31.1 \pm 2.7$	$715 \pm 35$	$27.7 \pm 3.0$ $30.6 \pm 3.5$	$70.8 \pm 4.2$	$27.2 \pm 3.8$ 297 + 43	
1110		StraTS	$80.8 \pm 0.4$	$42.4 \pm 0.9$	$80.4 \pm 0.7$	$41.8 \pm 1.8$	$79.5 \pm 1.6$	$40.2 \pm 3.1$	$78.7 \pm 1.9$	$38.4 \pm 4.2$	$78.4 \pm 2.0$	$37.5 \pm 4.4$	
1149		DuETT	$73.9 \pm 1.7$	$\frac{12.12}{35.8 \pm 2.3}$	$74.7 \pm 1.8$	$35.3 \pm 2.0$	$73.6 \pm 2.2$	$34.1 \pm 2.4$	$72.8 \pm 2.6$	$\frac{33.3 \pm 2.7}{33.3 \pm 2.7}$	$\frac{70.12}{72.3 \pm 2.7}$	$\frac{37.6 \pm 2.8}{32.6 \pm 2.8}$	
1150		Warpformer	75.9 ± 0.7	$37.3 \pm 2.2$	75.6 ± 0.8	$36.7 \pm 2.3$	73.8 ± 2.9	$34.3\pm4.1$	$72.8 \pm 3.4$	$33.0\pm4.6$	72.1 ± 3.7	$32.2 \pm 4.7$	
1150		Hi-Patch	<u>80.1 ± 1.1</u>	$43.3\pm2.2$	<u>79.7 ± 1.2</u>	$41.8\pm2.7$	<u>79.1 ± 1.5</u>	$40.2\pm3.4$	78.8 ± 1.5	$39.6\pm3.2$	78.7 ± 1.4	39.4 ± 3.0	
1151		GRU-D	88.5 ± 2.3	$54.6 \pm 3.7$	88.8 ± 2.1	$54.2 \pm 3.4$	88.0 ± 2.5	$50.4 \pm 7.5$	87.5 ± 2.8	$49.6 \pm 6.9$	86.4 ± 3.5	$47.2 \pm 8.6$	
1152		mTAND	79.6 ± 1.8	$28.6 \pm 1.9$	79.2 ± 1.9	$28.1 \pm 2.1$	$78.0 \pm 2.4$	$26.9 \pm 2.9$	$77.2 \pm 2.7$	$26.3 \pm 2.9$	$76.2 \pm 3.2$	$24.3 \pm 4.8$	
1150		DGM <sup>2</sup> -O	87.4 ± 0.6	$53.4 \pm 1.5$	87.3 ± 0.8	$53.2 \pm 1.7$	86.6 ± 1.6	$49.9 \pm 5.1$	$85.8 \pm 1.9$	47.7 ± 5.9	$85.2 \pm 2.2$	$45.7 \pm 6.7$	
1153	P19	MTGNN	$84.5 \pm 1.4$	$48.9 \pm 2.3$	84.8 ± 1.7	$49.8 \pm 3.1$	84.0 ± 1.9	$47.2 \pm 4.8$	83.3 ± 2.2	$45.5 \pm 5.5$	82.5 ± 2.9	$42.7 \pm 9.2$	
1154		Raindrop	$88.2 \pm 1.5$	$59.7 \pm 1.5$	88.1 ± 1.3	$59.8 \pm 1.4$	$87.8 \pm 1.2$	$59.1 \pm 1.7$	$87.6 \pm 1.1$	$58.5 \pm 1.9$	87.1 ± 1.5	$57.7 \pm 2.3$	
1154		StraTS	$90.6 \pm 0.9$	$56.4 \pm 3.0$	$91.0 \pm 0.9$	$56.3 \pm 2.3$	$91.0 \pm 0.9$	$56.0 \pm 2.4$	$90.8 \pm 1.0$	$55.1 \pm 3.0$	$90.4 \pm 1.3$	$54.4 \pm 3.3$	
1155		Wernformer	$85.2 \pm 1.0$	$55.7 \pm 1.0$ $55.2 \pm 5.6$	$84.8 \pm 1.1$ 012 ± 0.8	$53.9 \pm 0.8$ $55.1 \pm 5.6$	$84.7 \pm 1.0$ 01.4 ± 0.8	$53.3 \pm 1.0$ $56.0 \pm 4.8$	$84.3 \pm 1.4$ 015 ± 0.7	$52.7 \pm 2.1$ $56.4 \pm 4.2$	$84.4 \pm 1.3$	$52.5 \pm 2.0$ 56.2 ± 2.0	
1156		warptormer	<u>91.3 ± 0.8</u>	33.2 ± 3.0	<u>91.3 ± 0.8</u>	JJ.1 ± J.0	<u>91.4 ± 0.8</u>	J0.0 ± 4.8	<u>91.5 ± 0.7</u>	J0.4 ± 4.3	$\frac{91.2 \pm 0.8}{10.00}$	J0.2 ± 3.9	
1150		Hi-Patch	92.1 ± 0.4	60.7 ± 2.0	92.0 ± 0.4	60.6 ± 1.9	91.9 ± 0.5	$60.3 \pm 1.8$	91.9 ± 0.5	60.0 ± 1.7	91.6 ± 0.8	59.5 ± 1.9	
1157		GRU-D	70.0 ± 3.0	$32.1 \pm 4.1$	69.5 ± 2.6	$31.1 \pm 3.6$	69.2 ± 3.0	$31.0 \pm 4.4$	$68.3 \pm 3.6$	$30.1 \pm 5.3$	68.1 ± 3.7	$29.8 \pm 5.3$	
1150		mTAND	$80.5 \pm 2.1$	$42.8 \pm 4.0$	78.2 ± 3.4	$40.5 \pm 4.7$	$76.3 \pm 4.0$	37.7 ± 5.7	$75.6 \pm 3.9$	$36.6 \pm 5.6$	75.1 ± 3.9	$36.1 \pm 5.1$	
0011	DI INI	DGM <sup>2</sup> -O	$80.2 \pm 0.9$	$38.6 \pm 2.8$	$80.4 \pm 0.9$	$38.3 \pm 2.8$	$79.3 \pm 1.9$	$37.1 \pm 3.4$	$77.5 \pm 3.7$	$35.4 \pm 4.4$	$75.6 \pm 5.0$	$34.0 \pm 4.8$	
1159	PhysioNet	MIGNN	$68.9 \pm 4.1$	$25.8 \pm 4.8$	$69.3 \pm 4.3$	$26.6 \pm 4.5$	$69.0 \pm 4.8$	$26.3 \pm 5.2$	$68.3 \pm 5.2$	$25.4 \pm 4.8$	$67.2 \pm 5.4$	$24.4 \pm 4.8$	
		Kaindrop StroTS	$70.5 \pm 1.2$	$33.4 \pm 2.2$	$70.5 \pm 1.3$	$32.3 \pm 2.3$	$75.0 \pm 2.0$	$30.8 \pm 3.2$	$74.7 \pm 2.0$	$29.7 \pm 3.5$	$73.0 \pm 3.2$	$28.8 \pm 3.9$	
1160		DuFTT	$80.4 \pm 2.4$ 78 2 ± 2.8	$40.8 \pm 2.0$ $30.0 \pm 3.5$	$783 \pm 30$	$\frac{40.3 \pm 2.2}{30.0 \pm 3.7}$	$\frac{79.3 \pm 2.6}{767 \pm 3.7}$	$\frac{39.3 \pm 2.7}{37.0 \pm 4.5}$	$\frac{76.6 \pm 3.0}{75.0 \pm 3.8}$	$\frac{36.2 \pm 3.6}{37.0 \pm 4.6}$	$\frac{76.1 \pm 3.4}{74.0 \pm 4.3}$	$\frac{37.0 \pm 3.7}{35.0 \pm 5.0}$	
1161		Warnformer	782+10	333 + 21	777+16	$336 \pm 18$	758 + 34	318 + 30	$73.9 \pm 5.0$ $73.8 \pm 4.6$	302 + 41	$72.7 \pm 4.9$	292 + 42	
1101		Hi-Patch	812+34	$42.0 \pm 7.6$	806+34	$412 \pm 72$	708+34	$40.4 \pm 6.3$	703+33	30 8 + 5 8	787 + 34	30 3 + 5 5	
1162		III-I atch	01.2 ± 5.4	42.0 ± 7.0	00.0 ± 5.4	41.2 ± 7.2	17.0 ± 5.4	40.4 ± 0.5	17.5 ± 5.5	57.6 ± 5.6	70.7 ± 5.4	57.5 ± 5.5	
1163		GRU-D	$81.0 \pm 0.6$	$42.1 \pm 0.8$	$80.3 \pm 0.9$	$41.7 \pm 1.0$	$79.2 \pm 1.8$	$41.0 \pm 1.4$	$78.5 \pm 2.1$	$40.4 \pm 1.6$	$77.9 \pm 2.2$	$39.9 \pm 1.8$	
1105		mIAND	81.2 ± 0.2	$42.1 \pm 0.8$	80.4 ± 1.1	$41.9 \pm 1.2$	$79.7 \pm 1.4$	$41.0 \pm 1.7$	$79.3 \pm 1.4$	$40.4 \pm 2.0$	/8.8 ± 1.6	$39.8 \pm 2.3$	
1164	MIMIC-III	MTGNN	$78.8 \pm 0.5$	$34.2 \pm 0.9$ $34.5 \pm 1.4$	$78.3 \pm 0.8$	$33.9 \pm 1.1$ $24.0 \pm 1.2$	$77.0 \pm 1.2$	$33.4 \pm 1.2$ $23.5 \pm 1.5$	$77.3 \pm 1.3$ $76.3 \pm 2.5$	$33.1 \pm 1.2$ $22.8 \pm 1.0$	$75.6 \pm 1.5$	$32.0 \pm 1.4$ $32.2 \pm 2.4$	
1105	winwine-ini	Raindron	$78.0 \pm 1.1$	$34.3 \pm 1.4$ $33.7 \pm 0.0$	$77.5 \pm 1.0$	$34.0 \pm 1.3$ $33.5 \pm 0.0$	$77.1 \pm 2.2$ $76.4 \pm 2.1$	$33.3 \pm 1.3$ $32.8 \pm 1.4$	$76.3 \pm 2.3$	$32.0 \pm 1.9$ $32.5 \pm 1.4$	$75.0 \pm 3.2$	$32.2 \pm 2.4$ $32.3 \pm 1.4$	
1165		StraTS	$82.4 \pm 0.7$	$43.3 \pm 2.8$	82.1 + 0.6	$43.7 \pm 2.1$	81.7 + 0.8	$43.1 \pm 2.0$	$81.5 \pm 0.8$	$42.6 \pm 2.0$	$81.0 \pm 1.3$	$41.9 \pm 2.4$	
1166		DuETT	$78.0 \pm 0.5$	$\frac{10.0 \pm 0.0}{34.0 \pm 0.9}$	$77.2 \pm 1.0$	$\frac{33.7 \pm 0.8}{33.7 \pm 0.8}$	$76.6 \pm 1.2$	$33.3 \pm 1.0$	$76.4 \pm 1.2$	$33.0 \pm 1.0$	$76.1 \pm 1.3$	$32.6 \pm 1.3$	
		Warpformer	82.5 ± 0.5	$43.1 \pm 0.8$	81.7 ± 0.9	$42.5 \pm 1.2$	$81.2 \pm 1.1$	$42.1 \pm 1.2$	$80.6 \pm 1.5$	$41.8 \pm 1.3$	$80.0 \pm 1.9$	$41.3 \pm 1.6$	
1167		Hi-Patch	82.8 ± 0.3	44.3 ± 1.2	82.3 ± 0.6	44.1 ± 1.3	81.7 ± 1.0	43.6 ± 1.4	81.1 ± 1.5	$42.8 \pm 2.2$	80.5 ± 1.9	$42.0 \pm 2.6$	
1168													

## 1176 F.3 EFFICIENCY COMPARISON

We compared 8 leading IMTS models: Warpformer, GRU-D, GraFITi, tPatchGNN, CRU, mTAND, NeuralFlow, and Latent-ODE. All models are evaluated using the same batch size (32 for Human Activity, 128 for USHCN, 64 for PhysioNet, and 8 for MIMIC-III) to assess their training time per epoch and MSE. The results shown in Figure 6 indicate that the training time of our method ranks 5th on average across the four datasets. It is worth noting that our model employs a hierarchical structure to extract features at multiple scales within the IMTS, with each layer inevitably adding to the training time. In contrast, all the compared methods, except Warpformer, only extract features at the original scale, which is why our model is not the fastest. Nonetheless, our model's training time remains in the same order of magnitude as the fastest models. In this context, We believe that the trade-off of sacrificing some training time to extract richer features in IMTS is worthwhile.



Figure 6: Comparison of IMTS models in terms of efficiency: training time per epoch against error metric.





Figure 7: Visualization of views on three scales on Human Activity dataset.

