

Temporal Gaussian Copula For Clinical Multivariate Time Series Data Imputation

Ye Su

*Chongqing Institute of
Green and Intelligent Technology, CAS
University of Chinese Academy of Sciences*
Chongqing, China
suye241@mailsucas.ac.cn

Yuwen Chen*

*Chongqing Institute of
Green and Intelligent Technology, CAS*
Chongqing, China
chenyuwen@cigit.ac.cn

Hezhe Qiao

Singapore Management University
Singapore, Singapore
hezheqiao.2022@phdcs.smu.edu.sg

Di Wu

Southwest University
Chongqing, China
wudi.cigit@gmail.com

Lin Chen*

*Chongqing Institute of
Green and Intelligent Technology, CAS*
Chongqing, China
chenlin@cigit.ac.cn

Abstract—The imputation of the Multivariate time series (MTS) is particularly challenging since the MTS typically contains irregular patterns of missing values due to various factors such as instrument failures, interference from irrelevant data, and privacy regulations. Existing statistical methods and deep learning methods have shown promising results in time series imputation. In this paper, we propose a Temporal Gaussian Copula Model (TGC) for three-order MTS imputation. The key idea is to leverage the Gaussian Copula to explore the cross-variable and temporal relationships based on the latent Gaussian representation. Subsequently, we employ an Expectation-Maximization (EM) algorithm to improve robustness in managing data with varying missing rates. Comprehensive experiments were conducted on three real-world MTS datasets. The results demonstrate that our TGC substantially outperforms the state-of-the-art imputation methods. Additionally, the TGC model exhibits stronger robustness to the varying missing ratios in the test dataset. Our code is available at <https://github.com/chenlincigit/TGC-MTS>.

Index Terms—Multivariate Time Series Data, Gaussian Copula, Electronic Health Record, Temporal Missing Data Imputation.

I. INTRODUCTION

As a prevalent data in Electronic Health Records (EHR), clinical multivariate time series (MTS) data is widely used for various tasks for biomedical analysis, such as evaluating treatment effects [1], morbidity forecasts [2], and mortality projections [3], [4]. However, clinical MTS in the real world often contains many missing values due to various factors such as instrument failures, interference from irrelevant data, and privacy regulations [5], which presents a significant challenge to the application of existing machine learning methods and further impacts the accuracy of downstream analyses and

decision-making. Therefore, developing an accurate and effective strategy to handle missing data in MTS is essential for accurate clinical diagnosis.

While existing statistical and deep learning approaches [6]–[8] have shown promise in clinical MTS imputation, several challenges remain: (1) The complex relationships among clinical values at different time points are not fully captured. (2) Irregular and non-continuous recording intervals among patients can bias EHR data distribution, and current methods often struggle with varying sampling densities, resulting in suboptimal performance.

To tackle these limitations, we developed a cross-variable and longitudinal imputation approach, named Temporal Gaussian Copula (TGC), based on a Gaussian Copula operation to fully explore the complex relation between different clinical values and time-points. We utilize a conditional expectation formula based on EM optimization to ensure that all observed values are fully leveraged in imputing missing values, thereby accommodating any sampling density. Specifically, we first unfold the 3D clinical data as a single matrix by concatenating the clinical longitudinal information of the patients. Then, we employ the Gaussian copula to estimate complex multivariate distributions of variable-by-time correlations through the feature transformations of the latent Gaussian vectors. Finally, the model was solved using an effective approximate Expectation-Maximization (EM) algorithm, and the missing value was imputed using a defined inverse potential Gaussian vector transform. We conducted extensive experiments on three real-world medical MTS datasets at varying missing rates to evaluate the performance of TGC. The experimental results demonstrate that the TGC model exhibits the best average performance across different missing rates.

II. METHODOLOGY

The overview of TGC is illustrated in Fig 1. The TGC is primarily divided into four modules: (1) Three-order Tensor

*Lin Chen and Yuwen Chen are the corresponding authors. This research is supported in part by the National Nature Science Foundation of China (82030066 and 62371438) and also by the Chongqing Municipal Education Commission (HZ2021008, 24XJC630014, HZ2021017)

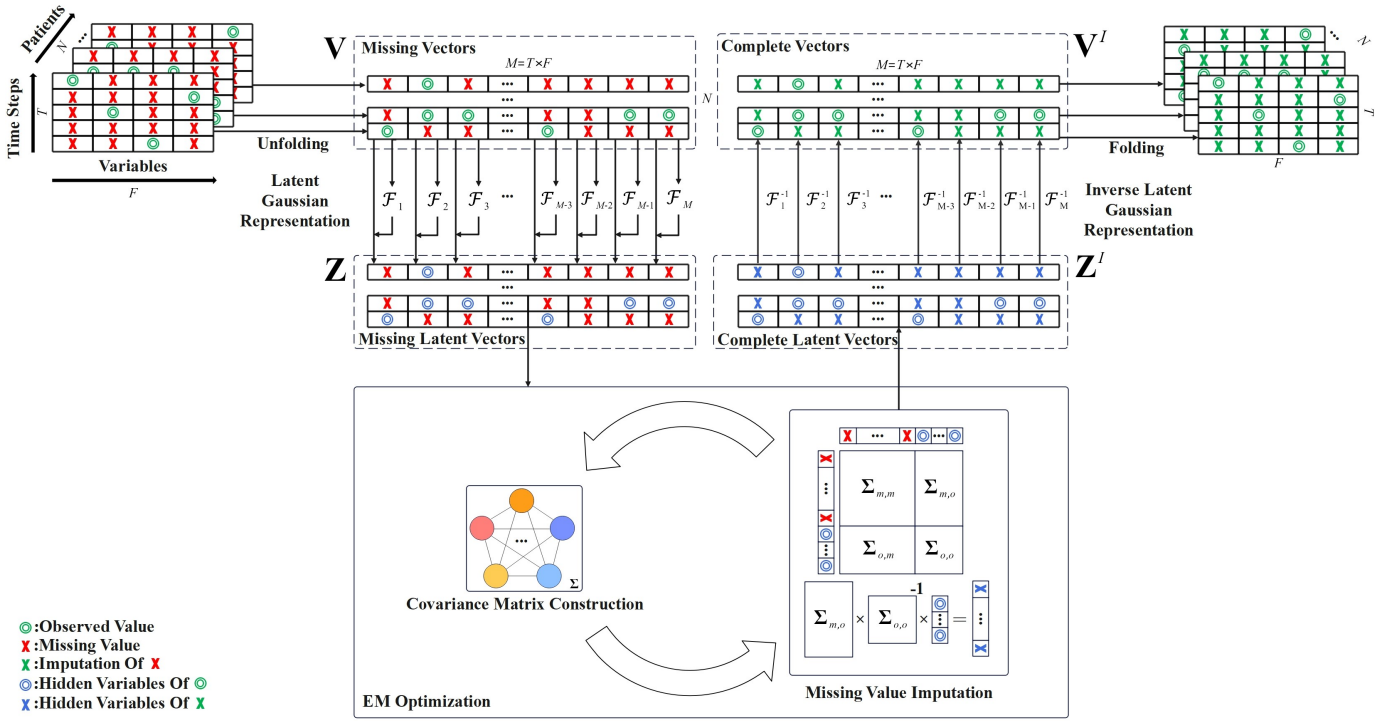


Fig. 1. The framework of TGC.

Unfolding. (2) Latent Gaussian Representation. (3) Covariance Matrix Construction. (4) Missing Value Estimation and Imputation. Subsequently, a detailed exposition of each constituent of the TGC will be presented.

A. Three-order Tensor Unfolding

We assume the MTS for a specific patient $\mathbf{X}_i \in \mathbb{R}^{T \times F}$, following the matrix Gaussian distribution with the parameters including mean matrix $\mathbf{M} \in \mathbb{R}^{T \times F}$, temporal covariance $\mathbf{U} \in \mathbb{R}^{T \times T}$, and covariance matrix $\mathbf{C} \in \mathbb{R}^{F \times F}$, which is formulated as the following:

$$\mathcal{P}(\mathbf{X}_i | \mathbf{M}, \mathbf{U}, \mathbf{C}) = (2\pi)^{-\frac{TF}{2}} |\mathbf{C}|^{-\frac{T}{2}} |\mathbf{U}|^{-\frac{F}{2}} e^{-\frac{1}{2} \text{Tr}[\mathbf{C}^{-1}(\mathbf{X}_i - \mathbf{M})^T \mathbf{U}^{-1}(\mathbf{X}_i - \mathbf{M})]}. \quad (1)$$

Thus, we can establish the temporal cross-variable correlations for missing data estimation by computing the covariance matrix Σ , using the matrices \mathbf{C} and \mathbf{U} .

However, estimating high-dimensional parameters of \mathbf{M} , \mathbf{C} , and \mathbf{U} is challenging due to the complexities inherent in matrix Gaussian. Therefore, to apply a simple yet effective missing data estimation approach, we can unfold the 3D tensor \mathcal{X} into a single 2D matrix $\mathbf{V} \in \mathbb{R}^{N \times M}$ which is formulated as:

$$\mathbf{V} = \{\text{vec}(\mathbf{X}_1), \text{vec}(\mathbf{X}_2), \dots, \text{vec}(\mathbf{X}_N)\}, \quad (2)$$

where $M = T \times F$, and $\text{vec}(\mathbf{X}_i)$ is the vectorization operator that transforms matrix \mathbf{X}_i into a column vector. Given each $\mathbf{v}_i \in \mathbb{R}^M$ follows the multivariate normal distribution, which is formulated as the following:

$$\mathcal{P}(\mathbf{v}_i | \mu, \Sigma) = (2\pi)^{-\frac{M}{2}} |\Sigma|^{-\frac{1}{2}} e^{-\frac{1}{2}(\mathbf{v}_i - \mu)^T \Sigma^{-1}(\mathbf{v}_i - \mu)}, \quad (3)$$

where μ is the mean vector and Σ is the corresponding covariance matrix. In this way, we only need to estimate μ and Σ using the multivariate normal distribution, which reduces the computational burden of estimating irrelevant parameters. After this transformation, the missing data imputation for MTS becomes that filling in the missing sample \mathbf{V}^m given the observed \mathbf{V}^o .

B. Latent Gaussian Representation

We then employ a mapping function $\mathcal{F}(\cdot)$ to transform the variables of \mathbf{v}_i into a latent Gaussian representation with Gaussian copula, which is formulated as:

$$z_{ij} = \mathcal{F}_j(v_{ij}), \quad (4)$$

where z_{ij} is the latent variable corresponding to v_{ij} . $\mathcal{F}_j(\cdot)$ represents the mapping function defined on the j -th variable in \mathbf{V} . Each $v_{ij} \in \mathbf{V}$ follows a multivariate Gaussian distribution $\mathcal{N}(\mu, \Sigma)$.

Based on the latent Gaussian representation \mathbf{z}_i , the prior probability distribution $\mathcal{P}(\mathbf{v}_i | \mu, \Sigma)$ can be flexibly represented by avoiding μ , which is formulated as:

$$\mathcal{P}(\mathbf{z}_i | \Sigma) = (2\pi)^{-\frac{M}{2}} |\Sigma|^{-\frac{1}{2}} e^{-\frac{1}{2} \mathbf{z}_i^T \Sigma^{-1} \mathbf{z}_i}, \quad (5)$$

where $\mathbf{z}_i = [z_{i1}, z_{i2}, \dots, z_{iM}] \in \mathbb{R}^M$ which can be obtained by the mapping \mathbf{v}_i . The mapping function $\mathcal{F}_j(\cdot)$ is implemented by:

$$\mathcal{F}_j = \Phi^{-1} \circ P_j, \quad (6)$$

where Φ is the standard normal cumulative distribution function (CDF). P_j is the CDF corresponding to the j -th variable.

The Φ is strictly monotonically increasing, making $\Phi^{-1}(x)$ exists, which is typically defined as the following:

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-\frac{t^2}{2}} dt. \quad (7)$$

The p_j manifests variations attributable to the diverse distributions of v_{ij} . For example, if v_{ij} follows a normal distribution $N(\mu, \sigma)$, it makes the $u_{ij} = P_j(\mathbf{v}_{ij})$ follows a uniform distribution on $[0, 1]$, which is formulated as the following:

$$P_j(x) = \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^x e^{-\frac{(t-\mu)^2}{2\sigma^2}} dt. \quad (8)$$

Similarly, through $z_{ij} = \Phi^{-1}(u_{ij})$, the uniform distribution can be transformed into a standard normal distribution $N(0, 1)$.

In order to map \mathbf{z}_{ij} to \mathbf{v}_{ij} back after imputing, we need define \mathcal{F}_j^{-1} for all features. To accommodate the presence of outliers, v_{ij} can take any value within the interval $[-\infty, +\infty]$. Therefore, for continuous variables, the probability density function is always greater than 0, and CDF is strictly monotonically increasing. By definition, the inverse of a strictly monotonically increasing function exists, hence P_j^{-1} is present.

However, if the variable $v_{ij} \in \{1, 2, \dots, K\}$, then CDF $P(v_{ij} \leq k + \gamma) = P(v_{ij} \leq k)$, $k \in \{1, 2, \dots, K\}$, $\gamma \in (0, 1)$, meaning it only satisfies monotonically increasing. Therefore, we need to define a specific P_j^{-1} to ensure the existence of the inverse mapping.

When considering a variable v_{ij} that belongs to the set $\{1, 2, \dots, K\}$, where its CDF is P_j , and $\Phi(z_{ij})$ falls within the interval $[P_j(k-1), P_j(k)]$ for $k \in \{1, 2, \dots, K\}$, we defines $\mathcal{F}_j^{-1}(z_{ij}) = P_j^{-1}(\Phi(z_{ij})) = k$.

So far, we define the hidden Gaussian vector mapping $\mathcal{F}_j(\cdot)$ and its inverse mapping $\mathcal{F}_j^{-1}(\cdot)$ for MTS data.

C. Covariance Matrix Construction

To fully explore the relation of latent variables, we leverage the maximum likelihood estimation (MLE) to estimate the correlation matrix Σ , which is formulated as the following:

$$\ell(\Sigma; \mathbf{Z}) = c - \frac{1}{2} \log(|\Sigma|) - \frac{1}{2} \text{Tr} \left[\Sigma^{-1} \frac{1}{N} \sum_{i=1}^N \mathbf{z}_i (\mathbf{z}_i)^\top \right], \quad (9)$$

where c is a constant, N denotes the number of patients.

When \mathbf{Z} represents complete data, by maximize Equation 9, we can derive the expression for Σ as follows:

$$\hat{\Sigma} = \frac{1}{N} \sum_{i=1}^N \mathbf{z}_i \mathbf{z}_i^\top. \quad (10)$$

It is worth noting that \mathbf{Z} contains missing values. We maximize the joint probability of between Σ and observed value \mathbf{Z}^o in the \mathbf{Z} .

$$\ell_o(\Sigma; \mathbf{Z}^o) = \frac{1}{N} \sum_{i=1}^N \int_{\mathbf{z} \in \mathbf{z}_i^o} \mathcal{P}(\mathbf{z}; \Sigma_{o,o}) dz, \quad (11)$$

where $\Sigma_{o,o}$ are submatrices of Σ , with rows and columns corresponding to $(\mathbf{z}_i^o, \mathbf{z}_i^o)$. The optimal $\hat{\Sigma}$ can be obtained by maximizing the joint probability.

$$\hat{\Sigma} = \arg \max_{\Sigma} \ell_o(\Sigma; \mathbf{Z}^o). \quad (12)$$

D. Missing Value Estimation

The missing values \mathbf{z}_i^m for the i -th patient can be obtained based on its observed values \mathbf{z}_i^o and covariance matrix Σ :

$$\mathbb{E}[\mathbf{z}_i^m | \mathbf{z}_i^o, \Sigma] = \Sigma_{m,o} \Sigma_{o,o}^{-1} \mathbf{z}_i^o, \quad (13)$$

where $\Sigma_{m,o}$ and $\Sigma_{o,o}$ are submatrices of Σ , with rows and columns corresponding to $(\mathbf{z}_i^m, \mathbf{z}_i^o)$ and $(\mathbf{z}_i^o, \mathbf{z}_i^o)$, respectively. We then utilize the inverse mapping $\mathcal{F}_j^{-1}(\cdot)$ to transform \mathbf{Z} back into \mathbf{V} .

E. EM Optimization

The EM algorithm has shown a remarkable ability to handle data with substantial missing values [9]. To make a more accurate estimation of Σ , we employ the EM algorithm to iteratively update the covariance matrix. Firstly, we initialize the missing part of the latent variable matrix \mathbf{Z}^m to zero and compute the corresponding covariance matrix $\Sigma^{(0)}$. The general procedure of the EM algorithm is as follows:

1) *E-Step*: Utilizing $\Sigma^{(t)}$ and \mathbf{z}_i^o , we recalculate \mathbf{z}_i^m through Equation 13 for all patients. We assume that the filled values are \mathbf{z}_i^I .

$$\mathbf{z}_i^I = \mathbb{E}[\mathbf{z}_i^m | \mathbf{z}_i^o, \Sigma^{(t)}]. \quad (14)$$

2) *M-Step*: We calculate the updated covariance matrix $\hat{\Sigma}^{(t+1)}$ using $G(\Sigma^{(t)}, \mathbf{z}_i^o)$. The function $G(\Sigma^{(t)}, \mathbf{z}_i^o)$ is related to Equation 10 and is expressed as follows:

$$G(\mathbf{z}_i^I, \mathbf{z}_i^o) = \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\mathbf{z}_i \mathbf{z}_i^\top | \mathbf{z}_i^o, \mathbf{z}_i^I]. \quad (15)$$

3) *Scale*: Given that Gaussian copulas correspond to the correlation matrix, we apply constraints on the iterative process of Σ through $P_{\mathcal{E}}$. The relationship is established as follows: $P_{\mathcal{E}} = \mathbf{D}^{-\frac{1}{2}} \Sigma \mathbf{D}^{-\frac{1}{2}}$, where \mathbf{D} is the diagonal matrix derived from Σ . This transformation ensures that the constraints are effectively incorporated into the iterative estimation of Σ .

III. EXPERIMENTS

A. Datasets Acquisition

The experiments are conducted on three public real-world healthcare datasets, including Medical Information Mart for Intensive Care (MIMIC) [10], PhysioNet 2012 Mortality Prediction Challenge (PhysioNet2012) [11] and PhysioNet in Cardiology Challenge 2019 (PhysioNet2019) [12]. The statistical results of the three datasets are shown in Table I.




Fig. 2. The performance of representative models at various levels of data missingness. The missing ratio on the horizontal axis represents the proportion of observations removed from the test set randomly, not the total missing rate of the test set.

TABLE I
GENERAL INFORMATION OF FOUR DATASETS USED IN THIS WORK

	MIMIC	PhysioNet-2012	PhysioNet-2019
#Total Samples	20,940	11,988	36,409
#Features	21	36	34
#Sequence Length	31	48	18
Original Missing Rate	68%	80%	79%

B. Implementation Detail

The data set is initially divided into the training set and the test set according to 80% and 20%. We randomly eliminate $E\%$ of observed values in the test set and use these values as ground truth to evaluate the imputation performance of models (See III-D). To fully leverage the computational power of statistical models designed for two-dimensional data, we reshaped the input data into a matrix format of size $N \times (TF)$. For deep learning-based models, we train the model on the training set and then perform interpolation on the masked test set to calculate the corresponding metrics. All the models are run through three rounds, with the average and variance of the corresponding metrics computed. We have carefully calibrated the parameters of these models based on the original papers and made necessary adjustments to adapt them to the varying characteristics of our datasets. Following the prior work [7], [8], MAE (mean absolute error), MRE (mean relative error), and RMSE (average square error of the radius) are used for evaluation. We implement these models using the PyPOTS and fancyimpute packages and employ the Adam optimizer for optimization on an NVIDIA GeForce RTX 4060 GPU.

C. Comparison with Competing Methods

We compare the TGC with the existing SOTA imputation algorithms by reporting the mean performance of the tested models across multiple trials where 20% to 80% of the observation data are randomly removed from the test set, which are shown in Table II. From Table II, we can see our TGC significantly outperforms other models, including statistical methods and deep learning-based methods on all databases.

TGC constructs a fully-connected covariance matrix to connect all the variables with all the time-points, thus it can learn complex dependencies and achieve the best results in MTS imputation.

D. Imputation With Varing Missing Ratios

To demonstrate the model’s effectiveness of handling the data with varying sampling densities, we further conduct experiments on the three datasets with different missing ratios, which is shown in Figure 2. In this experiment, all models were trained on datasets with the original missing rate (shown in Table I), and the testing data was randomly removed by 20% to 80% to simulate different missing levels. From the results, we observe that our TGC exhibits superior adaptability on the dataset with high missing ratios, indicating strong robustness. The statistical models perform poorly on large datasets with high missing rates, overall being inferior to deep learning models. Deep learning models are second only to TGC and higher than other models overall.

E. Ablation Study

In this subsection, to empirically evaluate the effectiveness of different modules of TGC, we have devised several variations of the TGC: (1) TGC-U, where the unfolding is replaced with a transformation of the form $(NT) \times F$, designed to differentiate it from the unfolding; (2) U-ISVD, which uses ISVD to fill in missing values after an unfolding operation.

The comparative analysis of the various models is summarized in Table III. These findings reveal that the TGC-U model failed to converge, emphasizing the crucial role of unfolding in assisting the model to grasp the complex interdependencies among different variables across various time points. Additionally, the U-ISVD model exhibits a considerably greater loss than the TGC model, thereby reinforcing the superior performance of the Gaussian copula.

IV. CONCLUSION

In this paper, we propose a TGC model for clinical MTS missing value estimation. We explore the cross-variable and temporal relationships based on the Gaussian latent representation. We employ an EM algorithm to iteratively refine the covariance matrix and impute missing values, thereby

TABLE II

THE COMPARISON OF TGC WITH OTHER METHODS. THE BEST PERFORMANCE PER DATASET IS HIGHLIGHTED IN BOLD. THE RESULTS IN BRACKETS REPRESENT THE AVERAGE STANDARD DEVIATION OF THREE EXPERIMENTS.

Model	MIMIC			PhysioNet2012			PhysioNet2019		
	MAE	MRE	RMSE	MAE	MRE	RMSE	MAE	MRE	RMSE
LOCF [13]	0.547 (0.003)	0.722 (0.004)	0.842 (0.003)	0.447 (0.001)	0.632 (0.003)	0.864 (0.043)	0.520 (0.001)	0.698 (0.002)	0.800 (0.002)
ISVD (2001) [14]	0.525 (0.002)	0.693 (0.002)	0.748 (0.003)	0.520 (0.002)	0.735 (0.002)	0.814 (0.031)	0.519 (0.002)	0.696 (0.001)	0.753 (0.007)
SI (2010) [15]	0.592 (0.001)	0.781 (0.001)	0.808 (0.002)	0.508 (0.001)	0.719 (0.001)	0.799 (0.034)	0.603 (0.002)	0.808 (0.000)	0.827 (0.006)
BRITS (2018) [16]	0.436 (0.002)	0.575 (0.002)	0.691 (0.003)	0.356 (0.001)	0.504 (0.001)	0.689 (0.038)	0.435 (0.002)	0.583 (0.002)	0.679 (0.006)
KNN (2019) [17]	0.683 (0.002)	0.901 (0.002)	0.951 (0.003)	0.639 (0.001)	0.903 (0.003)	1.008 (0.026)	0.743 (0.003)	0.996 (0.003)	1.037 (0.006)
GPVAE (2020) [18]	0.568 (0.004)	0.749 (0.004)	0.791 (0.004)	0.544 (0.015)	0.769 (0.020)	0.834 (0.032)	0.569 (0.004)	0.762 (0.005)	0.798 (0.008)
FiLM (2022) [19]	0.596 (0.002)	0.786 (0.002)	0.835 (0.004)	0.533 (0.002)	0.754 (0.001)	0.848 (0.031)	0.580 (0.002)	0.777 (0.001)	0.822 (0.006)
SCINet (2022) [6]	0.530 (0.002)	0.699 (0.003)	0.756 (0.004)	0.465 (0.002)	0.657 (0.003)	0.770 (0.034)	0.510 (0.004)	0.683 (0.004)	0.747 (0.008)
SAITS (2023) [7]	0.401 (0.004)	0.528 (0.005)	0.631 (0.005)	0.352 (0.002)	0.498 (0.003)	0.677 (0.042)	0.418 (0.002)	0.560 (0.002)	0.660 (0.006)
Koopa (2024) [8]	0.527 (0.005)	0.695 (0.006)	0.745 (0.005)	0.476 (0.007)	0.674 (0.011)	0.768 (0.035)	0.519 (0.007)	0.696 (0.009)	0.757 (0.009)
FreTS (2024) [20]	0.521 (0.008)	0.688 (0.011)	0.756 (0.009)	0.488 (0.007)	0.690 (0.009)	0.814 (0.043)	0.477 (0.008)	0.640 (0.011)	0.718 (0.010)
TGC(ours)	0.377 (0.001)	0.497 (0.002)	0.610 (0.002)	0.309 (0.001)	0.437 (0.002)	0.639 (0.042)	0.389 (0.001)	0.521 (0.001)	0.638 (0.006)

TABLE III

ABLATION STUDIES ON THE PROPOSED MODEL. THE BEST PERFORMANCE IS HIGHLIGHTED IN BOLD. THE RESULTS IN BRACKETS REPRESENT THE AVERAGE STANDARD DEVIATION OF THREE EXPERIMENTS.

Model	MIMIC			PhysioNet2012			PhysioNet2019		
	MAE	MRE	RMSE	MAE	MRE	RMSE	MAE	MRE	RMSE
TGC-U	0.760 (0.001)	1.002 (0.000)	1.003 (0.002)	0.709 (0.002)	1.003 (0.000)	1.01 (0.024)	0.746 (0.002)	1.000 (0.000)	0.995 (0.006)
U-ISVD	0.529 (0.002)	0.698 (0.003)	0.752 (0.002)	0.520 (0.002)	0.735 (0.002)	0.814 (0.031)	0.519 (0.002)	0.696 (0.001)	0.753 (0.007)
TGC	0.377 (0.001)	0.497 (0.002)	0.610 (0.002)	0.309 (0.001)	0.437 (0.002)	0.639 (0.042)	0.389 (0.001)	0.521 (0.001)	0.638 (0.006)

enhancing model adaptability. Extensive experiments validate the TGC model’s superiority in clinical data imputation. Moreover, TGC exhibits remarkable robustness in handling test datasets with high missing rates.

REFERENCES

- [1] R. Liu, K. M. Hunold, J. M. Caterino, and P. Zhang, “Estimating treatment effects for time-to-treatment antibiotic stewardship in sepsis,” *Nature machine intelligence*, vol. 5, no. 4, pp. 421–431, 2023.
- [2] Z. Nowroozilarki, A. Pakbin, J. Royalty, D. K. Lee, and B. J. Mortazavi, “Real-time mortality prediction using mimic-iv icu data via boosted nonparametric hazards,” in *2021 IEEE EMBS international conference on biomedical and health informatics (BHI)*. IEEE, 2021, pp. 1–4.
- [3] N. Hou, M. Li, L. He, B. Xie, L. Wang, R. Zhang, Y. Yu, X. Sun, Z. Pan, and K. Wang, “Predicting 30-days mortality for mimic-iii patients with sepsis-3: a machine learning approach using xgboost,” *Journal of translational medicine*, vol. 18, pp. 1–14, 2020.
- [4] T. Bergquist, T. Schaffter, Y. Yan, T. Yu, J. Prosser, J. Gao, G. Chen, E. Charzewski, Z. Nawalany, I. Brugere *et al.*, “Evaluation of crowd-sourced mortality prediction models as a framework for assessing artificial intelligence in medicine,” *Journal of the American Medical Informatics Association*, vol. 31, no. 1, pp. 35–44, 2024.
- [5] L. Ren, T. Wang, A. S. Seklouli, H. Zhang, and A. Bouras, “A review on missing values for main challenges and methods,” *Information Systems*, p. 102268, 2023.
- [6] M. Liu, A. Zeng, M. Chen, Z. Xu, Q. Lai, L. Ma, and Q. Xu, “Scinet: Time series modeling and forecasting with sample convolution and interaction,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 5816–5828, 2022.
- [7] W. Du, D. Côté, and Y. Liu, “Saits: Self-attention-based imputation for time series,” *Expert Systems with Applications*, vol. 219, p. 119619, 2023.
- [8] Y. Liu, C. Li, J. Wang, and M. Long, “Koopa: Learning non-stationary time series dynamics with koopman predictors,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [9] Y. Zhao and M. Udell, “Missing value imputation for mixed data via gaussian copula,” in *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, 2020, pp. 636–646.
- [10] A. E. Johnson, L. Bulgarelli, L. Shen, A. Gayles, A. Shammout, S. Horg, T. J. Pollard, S. Hao, B. Moody, B. Gow *et al.*, “Mimic-iv, a freely accessible electronic health record dataset,” *Scientific data*, vol. 10, no. 1, p. 1, 2023.
- [11] I. Silva, G. Moody, D. J. Scott, L. A. Celi, and R. G. Mark, “Predicting in-hospital mortality of icu patients: The physionet/computing in cardiology challenge 2012,” in *2012 computing in cardiology*. IEEE, 2012, pp. 245–248.
- [12] M. A. Reyna, C. S. Josef, R. Jeter, S. P. Shashikumar, M. B. Westover, S. Nemat, G. D. Clifford, and A. Sharma, “Early prediction of sepsis from clinical data: the physionet/computing in cardiology challenge 2019,” *Critical care medicine*, vol. 48, no. 2, pp. 210–217, 2020.
- [13] J. E. Overall, S. Tonidandel, and R. R. Starbuck, “Last-observation-carried-forward (lof) and tests for difference in mean rates of change in controlled repeated measurements designs with dropouts,” *Social Science Research*, vol. 38, no. 2, pp. 492–503, 2009.
- [14] O. Troyanskaya, M. Cantor, G. Sherlock, P. Brown, T. Hastie, R. Tibshirani, D. Botstein, and R. B. Altman, “Missing value estimation methods for dna microarrays,” *Bioinformatics*, vol. 17, no. 6, pp. 520–525, 2001.
- [15] R. Mazumder, T. Hastie, and R. Tibshirani, “Spectral regularization algorithms for learning large incomplete matrices,” *The Journal of Machine Learning Research*, vol. 11, pp. 2287–2322, 2010.
- [16] W. Cao, D. Wang, J. Li, H. Zhou, L. Li, and Y. Li, “Brits: Bidirectional recurrent imputation for time series,” *Advances in neural information processing systems*, vol. 31, 2018.
- [17] U. Pujianto, A. P. Wibawa, M. I. Akbar *et al.*, “K-nearest neighbor (k-nn) based missing data imputation,” in *2019 5th International Conference on Science in Information Technology (ICSITech)*. IEEE, 2019, pp. 83–88.
- [18] V. Fortuin, D. Baranchuk, G. Rätsch, and S. Mandt, “Gp-vae: Deep probabilistic time series imputation,” in *International conference on artificial intelligence and statistics*. PMLR, 2020, pp. 1651–1661.
- [19] T. Zhou, Z. Ma, Q. Wen, L. Sun, T. Yao, W. Yin, R. Jin *et al.*, “Film: Frequency improved legendre memory model for long-term time series forecasting,” *Advances in neural information processing systems*, vol. 35, pp. 12 677–12 690, 2022.
- [20] K. Yi, Q. Zhang, W. Fan, S. Wang, P. Wang, H. He, N. An, D. Lian, L. Cao, and Z. Niu, “Frequency-domain mlps are more effective learners in time series forecasting,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.