
SMART: Towards Pre-trained Missing-Aware Model for Patient Health Status Prediction

Zhihao Yu^{1,3} Xu Chu^{1,2,4,5*} Yujie Jin^{1,3} Yasha Wang^{3,4,5} Junfeng Zhao^{1,3,4,6}

¹School of Computer Science, Peking University

²Center on Frontiers of Computing Studies, Peking University, Beijing, China

³National Research and Engineering Center of Software Engineering, Peking University

⁴Key Laboratory of High Confidence Software Technologies, Ministry of Education, Beijing, China

⁵Peking University Information Technology Institute (Tianjin Binhai)

⁶Nanhu Laboratory, Jiaxing, China

yuzhihao@stu.pku.edu.cn, chu_xu@pku.edu.cn, wangyasha@pku.edu.cn

Abstract

Electronic health record (EHR) data has emerged as a valuable resource for analyzing patient health status. However, the prevalence of missing data in EHR poses significant challenges to existing methods, leading to spurious correlations and suboptimal predictions. While various imputation techniques have been developed to address this issue, they often obsess difficult-to-interpolate details and may introduce additional noise when making clinical predictions. To tackle this problem, we propose SMART, a Self-Supervised Missing-Aware Representation Learning approach for patient health status prediction, which encodes missing information via missing-aware temporal and variable attentions and learns to impute missing values through a novel self-supervised pre-training approach which reconstructs missing data representations in the latent space rather than in input space as usual. By adopting elaborated attentions and focusing on learning higher-order representations, SMART promotes better generalization and robustness to missing data. We validate the effectiveness of SMART through extensive experiments on six EHR tasks, demonstrating its superiority over state-of-the-art methods. Our code is available at <https://github.com/yzhHoward/SMART>.

1 Introduction

The rapid accumulation of electronic health record (EHR) data, driven by the widespread adoption of health information systems, has opened up new avenues for analyzing patient health status. EHR data in Intensive Care Units (ICUs) primarily captures patients' laboratory tests and vital signs, providing a rich resource for describing and analyzing their health conditions. These time series data have been leveraged to predict patient prognosis and physical status, attracting significant attention from both computer scientists and medical researchers. Various deep learning methods have been developed to exploit the patterns from EHR data [1, 2, 3, 4, 5, 6, 7, 8, 9], aiming to assist doctors in making informed decisions, improving work efficiency, and ultimately enhancing patient outcomes.

Recent advancements in EHR analysis can be divided into two categories. The first category focuses on improving the learning ability of time series by capturing feature correlations and exploring the collaborative relationships between variables with various techniques, such as convolutions [10], attentions [11, 12, 13, 14, 15], and graph neural networks [16, 17]. These methods have achieved promising results in patient health status prediction tasks, providing valuable insights into modeling the underlying patterns. However, a critical challenge is the prevalence of missing data in EHRs.

*Corresponding author.

Due to the fact that patients do not undergo all tests during each visit, EHR data is often highly sparse. These missing values can compromise the integrity of learned representations and potentially mislead models to capture spurious correlations, leading to erroneous predictions of patient health status. To address this issue, the other works [18, 19, 20] attempt to interpolate missing values in the input space by capturing the dynamics and regularizing the time series. Nevertheless, these methods may struggle with the complexity of the EHR data and concentrate on difficult-to-interpolate details instead of capturing the implicit semantic information. Furthermore, while these studies have demonstrated the potential of imputation methods to fill in missing values and enhance performance on clinical tasks in their settings, these approaches face limitations in maximizing the predictive power of the available data. By masking certain observations to serve as targets for imputation, a portion of the data is withheld from the model and cannot be fully utilized for predicting patient outcomes. Imperfect imputations can introduce additional noise into the data, causing unintended shifts in the data distribution and hindering the performance of subsequent tasks [21].

To tackle these challenges, we propose SMART, a Self-Supervised Missing-Aware RePresentTation Learning approach for patient health status prediction. The primary objective of our approach is to enable the model to encode missing information effectively. It performs variable-independent encoding on multivariate time series inputs and utilizes Missing-Aware RePresentTation learning blocks, namely MART blocks, to capture temporal and variable interactions, with both modules incorporating missing information. Instead of deploying the vanilla self-attention mechanism [22], we develop novel missing-aware attentions inside the MART block to cope with sparse data. Through stacking multiple blocks, the model can adequately learn correlations while perceiving missingness.

Another key innovation of our method lies in its two-stage training strategy. Inspired by previous work [18, 23, 24, 25] on imputing missing values, we adopt a self-supervised pre-training stage to enhance the ability to cope with sparse data, after which we fine-tune the model to accomplish the EHR tasks, such as mortality prediction. In contrast to previous works that perform reconstruction in the input space, our model conducts missing data reconstruction in the latent space. By randomly removing a portion of observations and training the model to reconstruct its representations, we enable the model to focus on learning higher-order representations which contain more semantic features rather than struggling with unnecessary details. This approach promotes better generalization and robustness to missing data. After pre-training, the embedding decoder for reconstruction is replaced with a task-specific decoder for patient health status prediction during fine-tuning. To further bridge the gap between these two tasks, we introduce a learnable vector that serves as the query for the proposed attention mechanism and as the basis for patient health status prediction.

We validate the effectiveness of SMART through extensive experiments on six different EHR tasks, including in-hospital mortality, sepsis, decompensation, phenotyping, and length of stay. Our results demonstrate that SMART, with its specific designs tailored to the missing characteristic of EHR data, significantly outperforms existing methods on all the metrics, setting a new state-of-the-art on these tasks. Furthermore, we showcase the robustness of our model under settings with higher missing rates, highlighting its potential for real-world applications in healthcare. Comparisons on the model efficiency illustrate that the proposed SMART is highly efficient, with lightweight parameters and fast training time among existing methods. Our work underscores the importance of integrating missing data awareness into representation learning for enhancing patient health status prediction and paves the way for more accurate and reliable clinical decision support systems.

2 Related Work

Clinical Predictive Models for EHR: Analyzing EHR data has become an increasingly popular research topic in the medical domain [15, 26, 27]. Numerous deep learning models have been developed to mine and leverage information from EHR data [28, 29, 30, 31, 32]. Early methods used recurrent networks and attentions to capture temporal information [11, 33, 34, 35]. For learning better representation and achieving higher performance, some works attempt to learn feature correlations via sophisticated network architectures [36, 37, 38]. For example, Baytas et al. [12] and Gao et al. [9] refine the design of recurrent networks by incorporating sampling intervals and disease progression in EHR data, thereby learning more comprehensive associations. Zhang et al. [39] design hierarchical recurrent encoders to capture both fragmented and global temporal variance. Ma et al. [13] and Ma et al. [15] adopt attention mechanisms to enhance biomarkers that have strong connections with outcomes. Other works try to incorporate medical knowledge from EHR data or

human expertise [6, 40, 41, 42]. For example, Zhang et al. [16] and Yu et al. [43] discover similar patients in the dataset and utilize their information to enhance learned representations and provide interpretations. Xu et al. [4] and Lu et al. [44] combine information from knowledge graphs for medical code data to improve predictions. Ye et al. [45] integrate medical knowledge graphs with patient disease progression paths to obtain better health representations. However, as discussed in the Introduction, all these methods face challenges in handling missing values. Most of them merely populate missingness with mean or front values during EHR data pre-processing, although it may be implausible and mislead the model to make wrong decisions. To address this challenge, we introduce missing-awareness mechanisms at multiple positions in SMART, enabling it to encode missingness and avoid the negative impact of missing values on representation learning.

Imputation Models for EHR Analysis: Imputation models are widely used in EHR analysis to handle missing values. Some works make efforts to interpolate missingness for learning better representations via clinical predictive tasks [21, 46, 47]. Specifically, Neil et al. [48], Che et al. [49], and Tan et al. [47] incorporate missing information into recurrent methods. Horn et al. [50] use a set-based approach and transform time series into sets of observations modeled by set functions insensitive to misalignment. Zhang et al. [51] consider sampling frequency and unify irregular time series in multiple scales. Nevertheless, these methods do not apply reconstruction losses and do not perform real imputations. The other works impute irregularly observed values to aligned reference points, notably recurrent methods [24], variational auto-encoders [23, 18], generative adversarial networks [24, 52], ordinary differential equations [19, 53], and probabilistic interpolation methods [20]. However, this paradigm randomly masks existing observations and conducts both clinical tasks and interpolation simultaneously, making already sparse data even sparser when predicting health status, leading to unsatisfactory performance compared to methods that do not mask observations. Recently, a pre-training approach [25] has been proposed for empowering models to impute missing values through self-supervised pre-training and then perform clinical tasks, allowing the model to execute EHR analyses on intact data while possessing imputation capabilities. Nonetheless, interpolating in the input space can lead to getting bogged down in optimizing some details rather than capturing the implicit information in the entire sequence, as well as skewing the underlying data distribution [21]. In contrast, SMART conducts missing data reconstruction in the latent space, enabling the model to focus on learning more semantic representations.

3 Methodology

In this section, we present the proposed SMART, a self-supervised missing-aware representation learning approach for predicting patients’ health status. In Figure 1, we give an overview of the information flow and highlight our designs.

3.1 Variable Independent Encoder

Notation: Let (x, m, y) denotes the EHR data of a specific patient, where the visit sequence $x \in \mathbb{R}^{T \times N}$ contains T visits with N variables. x_n^t records the value of the n -th indicator for the patient at visit t , accompanied by a binary mask m_n^t indicating whether the value is observed. For different patients, the visit length T can vary. To avoid the possible negative effects of varying intervals between visits, the observation intervals are aligned to hours following [13, 16, 17, 43]. Every patient corresponds to a label y , indicating the diagnosis or the outcome. Our goal is to improve the model’s representation learning capability to achieve better performance on y .

Given a patient’s EHR data, we apply a variable-independent encoding strategy to map (x, m) into a latent representation $h \in \mathbb{R}^{T \times N \times d}$, inspired by the success of previous works [15, 51, 54], where d is the dimension of latent space. Although variable-dependent approaches [9, 10] capture variable interactions and compress dimensions when embedding, such methods cannot learn further associations, i.e., the attention weight of each variable [13]. Unlike variable-independent recurrent models [13, 17, 43], we adopt linear projections to encode data, allowing the entire encoding process to be parallel and without accumulating noise from past visits. When handling variable values x and masks m , we combine them directly and let the model capture the interactions between them.

CLS Vector: When encoding the EHR data, we introduce a learnable vector which is concatenated before the time series. The role of this vector is to learn the information of the whole sequence in subsequent interactions and act as the pooled hidden state used for prediction (similar to the

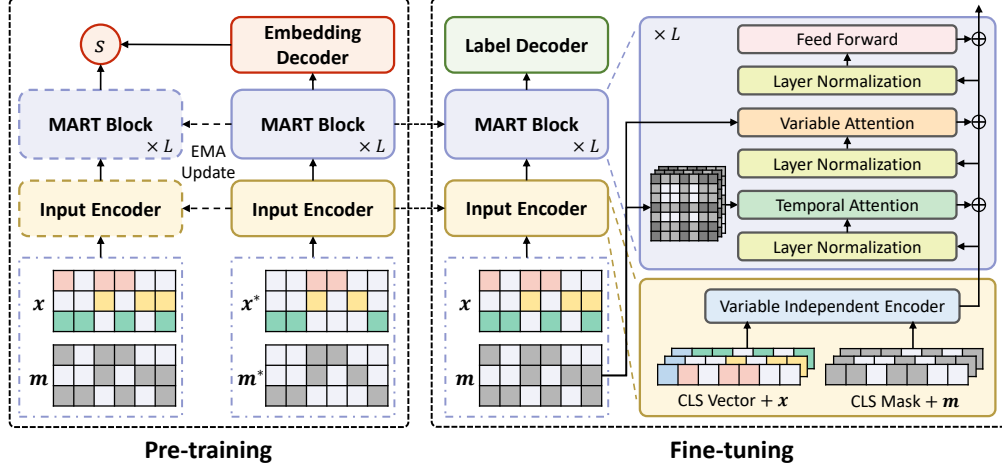


Figure 1: **Overview of SMART.** *Left:* Given EHR data with missingness, we randomly mask them on the existing observations and conduct reconstruction in the latent space. The reconstruction targets are generated by EMA updated parameters. *Right:* We illustrate the detailed architecture of the input encoder and the MART block. The input encoder embeds each variable (which can also be referred to as a biomarker) and missing mask into a separate hidden space. The MART block employs various techniques to capture feature interactions in both the temporal and variable dimensions while further encoding missing information.

[CLS] token in the language model such as BERT [55]). This design also bridges the gap between pre-training and fine-tuning tasks, as the information in the vector is used for both reconstruction and prediction. Specifically, we concatenate the CLS vector $v \in \mathbb{R}^{N \times d}$ at the temporal dimension before the hidden representation h , which can be formulated as $h_v = [v, h]$, where $h_v \in \mathbb{R}^{(T+1) \times N \times d}$. This process can be interpreted as concatenating a learnable parameter before the time series x and a True vector before the mask m , making the consultation sequence one step longer (as illustrated in Figure 1). Here, we concatenate on the embedding h instead of adding vectors to the input data since this procedure provides more flexibility in the learned representation space. At the end of this module, we introduce positional information to the representation h_v for subsequent interactions in the MART blocks using sinusoidal positional encoding [22].

3.2 MART Block

The MART block is the core module elaborated to learn the patient’s health representation. It is mainly composed of two attention mechanisms operating on the temporal and variable dimensions. To further mitigate the impact of missing data on representation learning, we introduce masks into the attention mechanism and strengthen the attention weights of existing observations.

Temporal Attention: The variable-independent encoder described above embeds patient information into representation $h_v \in \mathbb{R}^{(T+1) \times N \times d}$. In this section, we introduce how we calculate temporal weights leveraging mask information. Following the convention for self-attention mechanisms, we compute the query, key, and value via linear transformations of h_v for the temporal attention, denoted as Q_{temp} , K_{temp} , and V_{temp} . To incorporate mask information, we construct the temporal attention bias $B \in \mathbb{R}^{(T+1) \times (T+1) \times N}$. The bias $B_n^{i,j}$ between visit i and j for the n -th variable can be computed by

$$B_n^{i,j} = \begin{cases} 2, & m_n^i = \text{True} \text{ and } m_n^j = \text{True}, \\ 1 & m_n^i \text{ xor } m_n^j = \text{True}, \\ 0 & m_n^i = \text{False} \text{ and } m_n^j = \text{False}, \end{cases} \quad (1)$$

where $m_n^i = [\text{True}, m_n]$ and $[\cdot, \cdot]$ denotes concatenation of matrices. Then we obtain the temporal attention weights through $\text{softmax}\left(\frac{Q_{temp}K_{temp}^\top}{\sqrt{d}} + B\right)$, where both matrix multiplication and softmax are imposed on the temporal dimension. In this way, we strengthen information from

observed visits and suppress the others. At the same time, we do not completely block the missing visits, providing an opportunity for the model to interpolate them and utilize their information.

Variable Attention: We capture correlations between variables via the proposed variable attention. In contrast to previous approaches which calculate variable interactions separately for each visit [51, 54], we capture variable relationships from a global perspective of the patient. Given representation $\mathbf{h}_{temp} \in \mathbb{R}^{(T+1) \times N \times d}$ from the temporal attention, we get query \mathbf{Q}_{var} , key \mathbf{K}_{var} , and value \mathbf{V}_{var} for the variable attention mechanism as follows:

$$\begin{aligned} \mathbf{Q}_{var} &= \text{Linear}(\mathbf{h}_{temp}^0), \\ \mathbf{K}_{var} &= \text{Linear}\left(\sum_t \mathbf{h}_{temp}^t \text{ where } \mathbf{m}^t = \text{True}\right), \\ \mathbf{V}_{var} &= \text{Linear}(\mathbf{h}_{temp}). \end{aligned} \quad (2)$$

The query \mathbf{Q}_{var} is only obtained using the vector at the first step (\mathbf{h}_{temp}^0), which is also the location of learnable parameter \mathbf{v} inserted in the input encoder. This operation motivates these vectors to learn the overall health state of the patient. The key \mathbf{K}_{var} are acquired using the averaged embedding of all observed visits to minimize the effect of missingness in visits. Then we calculate the time-invariant correlations between the variables and get attention output by $\text{softmax}\left(\frac{\mathbf{Q}_{var}\mathbf{K}_{var}^\top}{\sqrt{d}}\right)\mathbf{V}_{var}$.

Between these attentions, we utilize layer normalizations [56] and skip connections [57] to avoid overfitting and accelerate convergence. Besides, a feed-forward constituted by linear projections and activation functions is used following the vanilla design of Transformer [22]. The MART block can be stacked in multiple layers to allow for sufficient interactions. The final health status of patients embedded by the MART blocks is denoted as $\mathbf{s} \in \mathbb{R}^{(T+1) \times N \times d}$.

3.3 Two-Stage Training Strategy

Pre-training Stage: To empower the model with missing imputation capabilities to enhance the learned representation, we propose a self-supervised pre-training method that occludes some of the observations and reconstructs their representation in the hidden space. Different from previous methods [25], we do not seek complex manual data augmentation, but simply remove some of the observations as the targets whose representations will be reconstructed. When generating the targets, we generate them randomly with a probability interval rather than with a fixed probability. This encourages the model to achieve better generalization across sequences with different sampling rates, rather than overfitting on missing probabilities. For the purpose of strengthening the reconstructing capability additionally and avoiding the model being trapped in the local minimum, we randomly sample the targets in each epoch of pre-training instead of using fixed data.

To emphasize a nontrivial pre-training task, we apply a self-motivated paradigm with an asymmetric architecture inspired by [58], as illustrated in Figure 1. Specifically, given the EHR data (\mathbf{x}, \mathbf{m}) , we randomly generate a mask $\hat{\mathbf{m}}$ to remove the existing observations partially and obtain augmented data $(\mathbf{x}^*, \mathbf{m}^*)$. Defining the modules being trained (input encoder, MART blocks, and embedding decoder) as f and the modules (input encoder and MART blocks) used to generate labels as \hat{f} , the reconstructions $\mathbf{s}_{pre-train}^*$ are generated by f using augmented data, while the reconstruction targets $\hat{\mathbf{s}}$ are acquired by embedding the original data via \hat{f} whose parameters are updated by exponential moving average (EMA) [59] of the parameters from f . The embedding decoder consists of an MLP with activation functions that uses the patient health representation $\mathbf{s}^* \in \mathbb{R}^{(T+1) \times N \times d}$ from augmented data as input and outputs the reconstructed health state $\mathbf{s}_{pre-train}^* \in \mathbb{R}^{(T+1) \times N \times d}$. This paradigm provides a smooth label update curve that avoids model oscillations and underfitted embedding decoder. The pre-training loss is computed by L_1 distance considering only the features of the removed data (i.e., the position where $\hat{\mathbf{m}} = \text{True}$) in the latent space:

$$\mathcal{L}_{pre-train} = \sum \hat{\mathbf{m}} \cdot \|\mathbf{s}_{pre-train}^* - \hat{\mathbf{s}}\|_1. \quad (3)$$

Since the model is trained to reconstruct the missing data, it is forced to learn the underlying structure of the data and the temporal relationships between variables, which can be beneficial for subsequent tasks. In particular, the CLS vector is not aligned in $\mathcal{L}_{pre-train}$ (i.e., the position of the CLS vector is

False in the mask \hat{m}), which allows the model to store the information of the whole sequence in the CLS vector and exploit it in the fine-tuning stage.

Fine-tuning Stage: After pre-training, we replace the embedding decoder with a task-specific decoder for patient health status prediction, such as classification task. While making predictions, we only use the representation at the first step (s^0), which is the position of the CLS vector. The label decoder is an MLP with layer normalizations and activations which can be simplified as a projection function $\mathbb{R}^{N \times d} \rightarrow \mathbb{R}^{|y|}$. In the fine-tuning stage, the parameters are updated by the task-specific loss (e.g. cross-entropy for classification). In the first few rounds of training, we freeze the parameters of the other modules and only update the label decoder to make the pre-trained parameters to be reserved. The proposed two-stage training procedure can be formulated in Algorithm 1:

Algorithm 1 Algorithm of SMART

- 1: **Input:** EHR data x and its mask m .
 - 2: **Output:** Patient health status prediction \hat{y} .
 - 3: **for all** E in pre-training epochs **do**
 - 4: Sample a mask \hat{m} with a probability interval p and generate augmented data (x^*, m^*)
 - 5: Generate reconstructions $s_{pre-train}^* \leftarrow f(x^*, m^*)$ and reconstruction targets $\hat{s} \leftarrow \hat{f}(x, m)$
 - 6: Update f by minimizing $\mathcal{L}_{pre-train}$
 - 7: Update $\hat{f} \leftarrow \text{EMA}(\hat{f}, f)$
 - 8: **end for**
 - 9: Freeze the parameters of the input encoder and the MART blocks
 - 10: **for all** E in fine-tuning epochs **do**
 - 11: **if** $E = \text{unfreeze epoch}$ **then**
 - 12: Unfreeze the parameters of the input encoder and the MART blocks
 - 13: **end if**
 - 14: Update parameters by minimizing the task-specific loss with prediction
 - 15: **end for**
-

4 Experiments

4.1 Experimental Settings

Datasets: We follow previous works [10, 13, 17, 51] to compare models on three EHR datasets, Cardiology [60, 61], Sepsis [62], and MIMIC-III [63]. The **Cardiology** dataset consists of 37 vital signs and biomarkers describing patients admitted to cardiac, medical, surgical, and trauma ICUs. Each record contains sparse measurements from the first 48 hours after admission. We follow the preprocessing procedures of previous works [13, 21] where the observation times are aligned to hours. After preprocessing, there are 11,988 patients and the observed rate is 24.7%. The prediction target is to determine in-hospital mortality. 13.8% of examples are in the positive class. The **Sepsis** dataset contains 34 vital signs and laboratory values relevant to sepsis. The EHR data are recorded once an hour including 40,335 patients from ICUs. The prediction task is to identify the patient’s risk of developing sepsis, with a positive rate of 7.3%. For convenience, we used only the records of the first 60 visits for each patient. The observed rate is only 19.8%. The **MIMIC-III** dataset is a multivariate time series dataset consisting of 17 physiological signals after pre-processing. It contains many clinical scenarios calling for accurate predictions over diversified clinical signals. We conduct four clinical tasks on the MIMIC-III including **in-hospital mortality**, **decompensation**, **phenotyping**, and **length of stay**. Additional details on these datasets can be found in Appendix A.1.

Evaluation Protocols: We assess the performance on the binary classification tasks (including Cardiology, Sepsis, in-hospital mortality, and decompensation) using the area under the precision-recall curve (AUPRC) and F1 Score. AUPRC is the most informative and primary evaluation metric when dealing with a highly imbalanced and skewed dataset [64] such as healthcare data. We calculate F1 Score focusing on positive patients, which is more relevant in clinical scenarios. Phenotyping is a multi-label classification, each label indicating the diagnosis of a phenotype. Therefore, we examine phenotyping on the macro and micro area under the receiver operating characteristic curve (AUROC), abbreviated as ma-ROC and mi-ROC, respectively. For length-of-stay prediction, we follow Harutyunyan et al. [65] to separate labels into multiple bins and evaluate the ma-ROC and

Table 1: Performance comparison with standard deviation on the six clinical tasks. The best results are marked in bold. The second best results are underlined.

Model	Cardiology		Sepsis		In-hospital Mortality	
	AUPRC(%)	F1 Score(%)	AUPRC(%)	F1 Score(%)	AUPRC(%)	F1 Score(%)
AdaCare	40.54±2.36	27.80±2.97	41.13±2.75	36.87±0.96	41.26±1.62	27.47±1.96
StageNet	47.01±2.94	38.27±2.74	59.11±4.63	54.89±3.28	51.77±3.62	<u>40.90±1.69</u>
ConCare	47.77±3.62	30.22±2.79	<u>74.83±3.47</u>	67.27±3.46	51.64±3.34	<u>35.76±2.58</u>
GRASP	47.93±5.73	31.90±2.60	74.47±3.79	66.75±3.41	<u>51.84±3.42</u>	38.00±1.75
SAFARI	45.58±3.59	11.06±2.09	72.06±1.74	69.06±1.14	<u>48.96±3.01</u>	33.96±3.83
PPN	<u>49.98±3.97</u>	37.64±3.19	72.70±4.99	68.08±2.71	51.24±2.96	40.60±1.83
RainDrop	<u>37.23±1.20</u>	23.85±3.39	72.82±3.68	67.58±2.54	46.21±2.71	36.29±3.20
Warpformer	48.10±3.95	<u>40.26±3.68</u>	73.03±2.53	65.93±1.79	51.49±1.89	36.83±3.12
PrimeNet	49.25±2.36	37.10±2.67	33.38±2.29	12.46±2.87	51.41±0.89	38.71±2.62
SMART	53.84±2.24	47.53±2.33	81.67±0.84	75.37±2.62	53.30±0.12	44.23±2.03

Model	Decompensation		Phenotyping		Length of Stay	
	AUPRC(%)	F1 Score(%)	ma-ROC(%)	mi-ROC(%)	ma-ROC(%)	mi-ROC(%)
AdaCare	52.41±1.76	51.62±0.84	63.28±0.62	73.75±0.26	66.06±0.97	82.11±0.40
StageNet	64.97±0.60	62.51±0.99	73.33±0.30	79.53±0.31	<u>69.54±0.39</u>	83.26±0.15
ConCare	65.50±1.30	59.73±2.25	71.05±0.36	78.03±0.21	<u>69.27±0.50</u>	83.19±0.12
GRASP	64.29±0.64	58.48±0.13	69.08±0.47	76.97±0.25	69.24±0.47	83.20±0.20
SAFARI	61.93±1.57	59.71±0.11	66.91±0.84	75.69±0.45	68.51±0.48	82.93±0.21
PPN	64.91±0.33	62.80±1.41	68.77±0.54	76.55±0.41	68.52±0.48	83.13±0.18
RainDrop	61.16±1.47	58.21±2.61	71.59±0.44	78.38±0.42	68.23±0.66	82.83±0.18
Warpformer	<u>68.21±1.64</u>	63.70±0.88	74.65±0.23	80.36±0.23	69.53±0.45	83.34±0.28
PrimeNet	<u>57.86±1.32</u>	51.59±1.71	68.90±0.17	<u>76.66±0.19</u>	67.88±0.40	82.89±0.18
SMART	71.26±0.81	67.38±0.67	76.13±0.14	81.41±0.18	70.29±0.44	83.92±0.24

mi-ROC. We randomly divide the data set into a training set containing 80% of the patients, a validation set of 10% patients, and a test set containing the remaining 10% instances. The model achieving the best AUPRC (or ma-ROC) on the validation set is evaluated on the test set. To eliminate the randomness, we conduct each experiment with three random seeds and report both the mean and standard deviation of the results. Results on more metrics are provided in Table 3. Hyperparameters and implementation details are introduced in Appendix A.2.

Baselines: We organize existing methods of handling EHR data into three paradigms. The first type utilizes various techniques to enhance representation learning of EHR data. Such as **AdaCare** [10], a multi-scale model focusing on extracting information from multiple levels, **StageNet** [9], which refines the design of LSTM by incorporating personalized disease stage development, and **ConCare** [13], which consists variable-independent encoders and self-attention mechanism to learn feature correlations. Works that employ global information or other patient information to aid in modeling also belong to this category, including **SAFARI** [17], which learns feature correlations from a group-wise perspective and integrates correlations by a graph neural network (GNN), **PPN** [43], a model that extracts typical patients and uses them help predicting and interpreting, and **GRASP** [16], which exploits similar patients and GNN to improve learned representation. The second paradigm is the state-of-the-art that delves into interpolating missingness and achieving better performance on clinical tasks, such as **RainDrop** [21], which utilizes GNN to capture the dependency among variables, **Warpformer** [51], which estimates the sampling frequency and interpolates missingness with attentions. The others are self-supervised pre-training methods, including **PrimeNet** [25], a model that is aware of missingness and adopts input reconstruction with contrastive learning to pre-train the model.

4.2 Experimental Results

4.2.1 Main Results

We include the overall comparison results on the six tasks in Table 1. SMART consistently outperforms existing baselines across all prediction metrics. Specifically, on the binary classification tasks, SMART achieves the results with an average improvement of 3.80% and 5.15% absolutely

Table 2: Ablation study of SMART on the Cardiology, Sepsis, and MIMIC-III in-hospital mortality.

SMART	Cardiology		Sepsis		In-hospital Mortality	
	AUPRC(%)	F1 Score(%)	AUPRC(%)	F1 Score(%)	AUPRC(%)	F1 Score(%)
Full	53.84±2.24	47.53±2.33	81.67±0.84	75.37±2.62	53.30±0.12	44.23±2.03
w/ Imputation	52.77±2.07	42.40±2.69	81.04±2.84	74.71±2.61	52.91±1.42	43.69±3.24
w/o Pre-training	52.26±3.14	46.60±2.07	79.54±3.22	74.26±2.65	51.84±0.95	41.69±2.95
w/o Mask	49.03±1.93	41.98±1.71	76.35±2.76	69.51±2.11	50.43±2.44	39.70±1.60
w/o Temporal Attention	49.28±2.94	40.53±2.31	65.22±1.44	58.47±2.36	46.64±1.63	33.37±3.82
w/o Variable Attention	52.27±2.53	44.75±2.82	80.47±2.36	74.65±2.39	50.98±0.62	41.70±2.92
w/o CLS Vector	52.96±0.34	46.42±1.77	77.56±2.89	71.53±2.56	50.78±1.64	43.86±2.85

on AUPRC and F1 Score compared to the best baseline, respectively. The results demonstrate the effectiveness of SMART in learning representations and predicting patient health status. AdaCare performs the worst in the baseline, which may be due to the fact that its convolutional structure only averages the proximity visits and does not perceive the missing ones. Interestingly, although StageNet and PrimeNet show competitive performance on some tasks such as Cardiology and in-hospital mortality, they perform surprisingly poorly on the Sepsis dataset. This is because the Sepsis dataset has a lower observed rate, which makes it more challenging for models to learn effective representations. Some recurrent models, including ConCare, GRASP, and PPN, exhibit robust performance across all datasets among the baselines. Nevertheless, with the capability of encoding missingness, SMART achieves remarkable performance improvements. As we observed, on the datasets with higher missing rate (Cardiology and Sepsis), SMART outperforms the other methods by a larger margin, indicating its robustness to missing values. On the phenotyping and length-of-stay prediction, our proposed model achieves the best results as well, showing its generalization ability to different clinical scenarios. The results also show that the performance of SMART is more stable than other methods, as indicated by the smaller standard deviation.

4.2.2 Ablation Study

To evaluate the effectiveness of each component in SMART, we conduct an ablation study on the Cardiology, Sepsis, and in-hospital mortality. Firstly, we investigate the effects of different pre-training strategies on model performance. We introduce two variants of the self-supervised pre-training strategy: (1) imputing the missing values in the input space like previous methods [18, 24, 25] (**w/ Imputation**), and (2) directly training the model for patient health status prediction without the proposed self-supervised pre-training (**w/o Pre-training**). As the results shown in Table 2, we observe that although combining the pre-training strategy in the input space can improve the performance compared to models without pre-training, it is still inferior to the proposed strategy that reconstructs in the representation space. These findings highlight the necessity of the pre-training strategy that integrates the missing imputation ability by reconstructing latent representations.

Additionally, we explore the importance of different components and designs in SMART, and the results are also presented in Table 2. In particular, we compare with the following reduced variants: (1) **w/o Mask**, which removes the mask information totally in SMART, including mask in the input encoder and the attentions in the MART blocks; (2) **w/o Temporal Attention**, which removes the temporal attention mechanism in SMART; (3) **w/o Variable Attention**, which removes the variable attention mechanism in SMART; and (4) **w/o CLS Vector**, which removes the CLS vector in the input encoder and using the representation at last observation as query in the variable attention and prediction. We observe that all components significantly contribute to the improvement. Notably, incorporating missing information is important for the model to learn high-quality representations. The results show that the temporal and variable attention mechanisms are crucial and fundamental in capturing temporal and variable dependencies, respectively. The CLS vector also plays a critical role in improving the effect of pre-training since it narrows the difference between the two training stages. By considering these components together, SMART is able to capture the intricate temporal relationships and characteristics inner sparse EHR data.




Figure 2: Performance on different observed ratio of EHR.




Figure 3: Training time, parameters, and AUPRC(%) of all models on the three datasets. The size of the circle represents the number of parameters. The GPU runtime is counted at the start of training.

4.2.3 Effect of Missingness

To further investigate the impact of missingness in the data on performance, we conduct a comprehensive experiment on the Cardiology, Sepsis, and in-hospital mortality tasks by varying the observed rate from 10% to 100%. The results of AUPRC are shown in Figure 2. AdaCare shows the most significant performance degradation as the observed rate decreases, which confirms that the convolutional architecture inside AdaCare is sensitive to missing values. Though Warpformer is stable on the Sepsis, it cannot handle the missingness well on the Cardiology and in-hospital mortality. Among the baselines, StageNet and PPN exhibit robustness to missing values. However, there is still a large margin compared to SMART. Especially, on the Sepsis, we observe only a very small decline on SMART despite only 10% of the available data. SMART performs outstandingly under different missing rate scenarios, demonstrating its fruitful results in missingness perception.

4.3 Model Efficiency

We conduct runtime and parameter comparisons between our method and baselines on the Cardiology, Sepsis, and in-hospital mortality prediction tasks. The number of parameters, the averaged runtime (minutes) on GPU (counted at the start of training), and the AUPRC of different methods are reported in Figure 3 for comparisons. Since the Sepsis dataset contains a larger number of patients, the model trained on it takes longer on average. As illustrated, RainDrop has the largest number of parameters and its training time is several times that of some of the models, such as AdaCare, ConCare, GRASP, Warpformer, and SMART, yet its prediction performance is not satisfactory and unstable on different datasets. StageNet, PPN, and PrimeNet achieve competitive performance on some of the datasets, however their training times are very long. Although AdaCare has the fewest parameters and the shortest training time, it also has one of the worst performances. In a nutshell, SMART is lightweight, fast, and reliable among all the models, reaching the best performance with only fewer parameters and shorter training time, which confirms its efficiency and effectiveness. Nevertheless, it is worth

mentioning that due to the temporal attention mechanism in SMART performs in the visit dimension, its training time may grow quadratically as the number of visits increases.

5 Conclusions and Limitations

In this work, we propose SMART, a self-supervised pre-trained model that is designed to handle the challenges of missingness and irregularity in EHR data. We introduce a two-stage training strategy that integrates imputing missingness in the representation space to enhance performance on clinical tasks. We elaborate on a novel MART block that captures temporal and variable interactions by introducing masks into the attention mechanism. Extensive experiments on six EHR tasks demonstrate that SMART outperforms existing baselines. The ablation study shows that all components in SMART contribute to the improvement. We further display that SMART is lightweight, efficient, and robust to missing values and achieves stable performance across different missing rates. In the future, we plan to provide more insights into the decision-making process and make it more explainable.

Despite the promising results obtained in our work, it is important to acknowledge its limitations. Like all models with temporal attention, the time and space complexity of our model is quadratic, which means that the model may face slow speed or insufficient memory when the input sequence is too long. Besides, we evaluated the experimental results with three datasets in reference to previous work, while there remain other datasets that could be taken into account.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (No.82241052).

References

- [1] Huan Song, Deepta Rajan, Jayaraman Thiagarajan, and Andreas Spanias. Attend and diagnose: Clinical time series analysis using attention models. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.
- [2] Junyu Luo, Muchao Ye, Cao Xiao, and Fenglong Ma. Hitanet: Hierarchical time-aware attention networks for risk prediction on electronic health records. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 647–656, 2020.
- [3] Liantao Ma, Xinyu Ma, Junyi Gao, Xianfeng Jiao, Zhihao Yu, Chaohe Zhang, Wenjie Ruan, Yasha Wang, Wen Tang, and Jiangtao Wang. Distilling knowledge from publicly available online emr data to emerging epidemic for prognosis. In *Proceedings of the Web Conference 2021*, pages 3558–3568, 2021.
- [4] Yongxin Xu, Kai Yang, Chaohe Zhang, Peinie Zou, Zhiyuan Wang, Hongxin Ding, Junfeng Zhao, Yasha Wang, and Bing Xie. Vecocare: visit sequences-clinical notes joint learning for diagnosis prediction in healthcare data. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, pages 4921–4929, 2023.
- [5] Yongxin Xu, Xu Chu, Kai Yang, Zhiyuan Wang, Peinie Zou, Hongxin Ding, Junfeng Zhao, Yasha Wang, and Bing Xie. Seqcare: Sequential training with external medical knowledge graph for diagnosis prediction in healthcare data. In *Proceedings of the ACM Web Conference 2023*, pages 2819–2830, 2023.
- [6] Pengcheng Jiang, Cao Xiao, Adam Richard Cross, and Jimeng Sun. Graphcare: Enhancing healthcare predictions with personalized knowledge graphs. In *The Twelfth International Conference on Learning Representations*, 2024.
- [7] Isotta Landi, Benjamin S Glicksberg, Hao-Chih Lee, Sarah Cherng, Giulia Landi, Matteo Danieletto, Joel T Dudley, Cesare Furlanello, and Riccardo Miotto. Deep representation learning of electronic health records to unlock patient stratification at scale. *NPJ digital medicine*, 3(1):96, 2020.
- [8] Riccardo Miotto, Li Li, Brian A Kidd, and Joel T Dudley. Deep patient: an unsupervised representation to predict the future of patients from the electronic health records. *Scientific reports*, 6(1):1–10, 2016.
- [9] Junyi Gao, Cao Xiao, Yasha Wang, Wen Tang, Lucas M Glass, and Jimeng Sun. Stagenet: Stage-aware neural networks for health risk prediction. In *Proceedings of The Web Conference 2020*, pages 530–540, 2020.

- [10] 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 scale-adaptive feature extraction and recalibration. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 825–832, 2020.
- [11] Edward Choi, Mohammad Taha Bahadori, Jimeng Sun, Joshua Kulas, Andy Schuetz, and Walter Stewart. Retain: An interpretable predictive model for healthcare using reverse time attention mechanism. *Advances in neural information processing systems*, 29, 2016.
- [12] Inci M Baytas, Cao Xiao, Xi Zhang, Fei Wang, Anil K Jain, and Jiayu Zhou. Patient subtyping via time-aware lstm networks. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 65–74, 2017.
- [13] Liantao Ma, Chaohe Zhang, Yasha Wang, Wenjie Ruan, Jiangtao Wang, Wen Tang, Xinyu Ma, Xin Gao, and Junyi Gao. Concare: Personalized clinical feature embedding via capturing the healthcare context. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 833–840, 2020.
- [14] Tian Bai, Shanshan Zhang, Brian L Egleston, and Slobodan Vucetic. Interpretable representation learning for healthcare via capturing disease progression through time. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 43–51, 2018.
- [15] Liantao Ma, Chaohe Zhang, Junyi Gao, Xianfeng Jiao, Zhihao Yu, Yinghao Zhu, Tianlong Wang, Xinyu Ma, Yasha Wang, Wen Tang, et al. Mortality prediction with adaptive feature importance recalibration for peritoneal dialysis patients. *Patterns*, 4(12), 2023.
- [16] Chaohe Zhang, Xin Gao, Liantao Ma, Yasha Wang, Jiangtao Wang, and Wen Tang. Grasp: generic framework for health status representation learning based on incorporating knowledge from similar patients. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 715–723, 2021.
- [17] Xinyu Ma, Yasha Wang, Xu Chu, Liantao Ma, Wen Tang, Junfeng Zhao, Ye Yuan, and Guoren Wang. Patient health representation learning via correlational sparse prior of medical features. *IEEE Transactions on Knowledge and Data Engineering*, 2022.
- [18] Satya Narayan Shukla and Benjamin Marlin. Multi-time attention networks for irregularly sampled time series. In *International Conference on Learning Representations*, 2021.
- [19] 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.
- [20] SeungHyun Kim, Hyunsu Kim, Eunggu Yun, Hwangrae Lee, Jaehun Lee, and Juho Lee. Probabilistic imputation for time-series classification with missing data. In *International Conference on Machine Learning*, pages 16654–16667. PMLR, 2023.
- [21] 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.
- [22] 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.
- [23] Satya Narayan Shukla and Benjamin Marlin. Interpolation-prediction networks for irregularly sampled time series. In *International Conference on Learning Representations*, 2019.
- [24] Xiaoye Miao, Yangyang Wu, Jun Wang, Yunjun Gao, Xudong Mao, and Jianwei Yin. Generative semi-supervised learning for multivariate time series imputation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 8983–8991, 2021.
- [25] Ranak Roy Chowdhury, Jiacheng Li, Xiyuan Zhang, Dezhi Hong, Rajesh K Gupta, and Jingbo Shang. Primenet: Pre-training for irregular multivariate time series. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 7184–7192, 2023.
- [26] Xiongcai Cai, Oscar Perez-Concha, Enrico Coiera, Fernando Martin-Sanchez, Richard Day, David Roffe, and Blanca Gallego. Real-time prediction of mortality, readmission, and length of stay using electronic health record data. *Journal of the American Medical Informatics Association*, 23(3):553–561, 2016.

- [27] Xinlu Zhang, Shiyang Li, Zhiyu Chen, Xifeng Yan, and Linda Ruth Petzold. Improving medical predictions by irregular multimodal electronic health records modeling. In *International Conference on Machine Learning*, pages 41300–41313. PMLR, 2023.
- [28] Tengfei Ma, Cao Xiao, and Fei Wang. Health-atm: A deep architecture for multifaceted patient health record representation and risk prediction. In *Proceedings of the 2018 SIAM International Conference on Data Mining*, pages 261–269. SIAM, 2018.
- [29] Yuqi Si, Jingcheng Du, Zhao Li, Xiaoqian Jiang, Timothy Miller, Fei Wang, W Jim Zheng, and Kirk Roberts. Deep representation learning of patient data from electronic health records (ehr): A systematic review. *Journal of biomedical informatics*, 115:103671, 2021.
- [30] Chaohe Zhang, Xu Chu, Liantao Ma, Yinghao Zhu, Yasha Wang, Jiangtao Wang, and Junfeng Zhao. M3care: Learning with missing modalities in multimodal healthcare data. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 2418–2428, 2022.
- [31] Jingyue Gao, Xiting Wang, Yasha Wang, Zhao Yang, Junyi Gao, Jiangtao Wang, Wen Tang, and Xing Xie. Camp: Co-attention memory networks for diagnosis prediction in healthcare. In *2019 IEEE international conference on data mining (ICDM)*, pages 1036–1041. IEEE, 2019.
- [32] Tian Bai and Slobodan Vucetic. Improving medical code prediction from clinical text via incorporating online knowledge sources. In *The World Wide Web Conference*, pages 72–82, 2019.
- [33] Xiaohan Li, Shu Wu, and Liang Wang. Blood pressure prediction via recurrent models with contextual layer. In *Proceedings of the 26th International Conference on World Wide Web*, pages 685–693, 2017.
- [34] Yanbo Xu, Siddharth Biswal, Shripasad R Deshpande, Kevin O Maher, and Jimeng Sun. Raim: Recurrent attentive and intensive model of multimodal patient monitoring data. In *Proceedings of the 24th ACM SIGKDD international conference on Knowledge Discovery & Data Mining*, pages 2565–2573, 2018.
- [35] Fenglong Ma, Radha Chitta, Jing Zhou, Quanzeng You, Tong Sun, and Jing Gao. Dipole: Diagnosis prediction in healthcare via attention-based bidirectional recurrent neural networks. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1903–1911, 2017.
- [36] Chengxi Zang and Fei Wang. Scehr: Supervised contrastive learning for clinical risk prediction using electronic health records. In *Proceedings. IEEE International Conference on Data Mining*, volume 2021, page 857. NIH Public Access, 2021.
- [37] Yujie Feng, Jiangtao Wang, Yasha Wang, and Sumi Helal. Completing missing prevalence rates for multiple chronic diseases by jointly leveraging both intra- and inter-disease population health data correlations. In *Proceedings of the Web Conference 2021*, pages 183–193, 2021.
- [38] Chang Lu, Tian Han, and Yue Ning. Context-aware health event prediction via transition functions on dynamic disease graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 4567–4574, 2022.
- [39] Jinghe Zhang, Kamran Kowsari, James H Harrison, Jennifer M Lobo, and Laura E Barnes. Patient2vec: A personalized interpretable deep representation of the longitudinal electronic health record. *IEEE Access*, 6:65333–65346, 2018.
- [40] Edward Choi, Mohammad Taha Bahadori, Le Song, Walter F Stewart, and Jimeng Sun. Gram: graph-based attention model for healthcare representation learning. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 787–795, 2017.
- [41] Fenglong Ma, Quanzeng You, Houping Xiao, Radha Chitta, Jing Zhou, and Jing Gao. Kame: Knowledge-based attention model for diagnosis prediction in healthcare. In *Proceedings of the 27th ACM international conference on information and knowledge management*, pages 743–752, 2018.
- [42] Junyi Gao, Chaoqi Yang, Joerg Heintz, Scott Barrows, Elise Albers, Mary Stapel, Sara Warfield, Adam Cross, and Jimeng Sun. Medml: fusing medical knowledge and machine learning models for early pediatric covid-19 hospitalization and severity prediction. *Iscience*, 25(9), 2022.
- [43] Zhihao Yu, Chaohe Zhang, Yasha Wang, Wen Tang, Jiangtao Wang, and Liantao Ma. Predict and interpret health risk using ehr through typical patients. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1506–1510. IEEE, 2024.
- [44] Chang Lu, Chandan K Reddy, Prithwish Chakraborty, Samantha Kleinberg, and Yue Ning. Collaborative graph learning with auxiliary text for temporal event prediction in healthcare. 2021.

- [45] Muchao Ye, Suhan Cui, Yaqing Wang, Junyu Luo, Cao Xiao, and Fenglong Ma. Medpath: Augmenting health risk prediction via medical knowledge paths. In *Proceedings of the Web Conference 2021*, pages 1397–1409, 2021.
- [46] Qianli Ma, Sen Li, and Garrison W Cottrell. Adversarial joint-learning recurrent neural network for incomplete time series classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(4):1765–1776, 2020.
- [47] Qingxiong Tan, Mang Ye, Baoyao Yang, Siqi Liu, Andy Jinhua Ma, Terry Cheuk-Fung Yip, Grace Lai-Hung Wong, and PongChi Yuen. Data-gru: Dual-attention time-aware gated recurrent unit for irregular multivariate time series. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 930–937, 2020.
- [48] Daniel Neil, Michael Pfeiffer, and Shih-Chii Liu. Phased lstm: Accelerating recurrent network training for long or event-based sequences. *Advances in neural information processing systems*, 29, 2016.
- [49] 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.
- [50] Max Horn, Michael Moor, Christian Bock, Bastian Rieck, and Karsten Borgwardt. Set functions for time series. In *International Conference on Machine Learning*, pages 4353–4363. PMLR, 2020.
- [51] 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*, pages 3273–3285, 2023.
- [52] Ying Zhang, Baohang Zhou, Xiangrui Cai, Wenya Guo, Xiaoke Ding, and Xiaojie Yuan. Missing value imputation in multivariate time series with end-to-end generative adversarial networks. *Information Sciences*, 551:67–82, 2021.
- [53] 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, 2024.
- [54] Yuqi Nie, Nam H Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A time series is worth 64 words: Long-term forecasting with transformers. In *The Eleventh International Conference on Learning Representations*, 2023.
- [55] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186, 2019.
- [56] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [57] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [58] Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15619–15629, 2023.
- [59] J Stuart Hunter. The exponentially weighted moving average. *Journal of quality technology*, 18(4):203–210, 1986.
- [60] 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, physiotookit, and physionet: components of a new research resource for complex physiologic signals. *circulation*, 101(23):e215–e220, 2000.
- [61] 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 Computing in Cardiology*, pages 245–248. IEEE, 2012.
- [62] Matthew A Reyna, Christopher S Josef, Russell Jeter, Supreeth P Shashikumar, M Brandon Westover, 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.

- [63] 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.
- [64] Jesse Davis and Mark Goadrich. The relationship between precision-recall and roc curves. In *Proceedings of the 23rd international conference on Machine learning*, pages 233–240, 2006.
- [65] 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.
- [66] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

A Additional Experimental Results

A.1 Details of Datasets

When preprocessing all the datasets, we conduct z-score normalizations (subtract the mean and divide the deviation of the observed values) on the continuous features and one-hot encoding on the categorical features. We align the observation times to hours and pad the missing values with zeros. We use the first 48 hours of data for the Cardiology and the MIMIC-III mortality prediction dataset and the first 60 visits for the Sepsis dataset. We refer to the procedures from Harutyunyan et al. [65] when generating samples from the MIMIC-III dataset. We split the data into training, validation, and test sets with a ratio of 8:1:1. The Cardiology dataset can be obtained at <https://physionet.org/content/challenge-2012/1.0.0/>. The Sepsis dataset can be obtained at <https://physionet.org/content/challenge-2019/1.0.0/>. The MIMIC-III dataset can be obtained at <https://physionet.org/content/mimiciii/1.4/>. Please follow the instructions on the PhysioNet website to access the data. We provide more details about the tasks on the MIMIC-III dataset as follows:

In-hospital Mortality: The goal of this task is to predict whether a patient will die at the end of hospitalization using the first 48 hours of data. We extract 21,139 records each containing 17 physiological variables with 43.3% observed rate. There are 13.2% of patients who have positive labels. In this task, the model can only observe the first 48 hours, which is advantageous from a practical point of view, as the earlier the clinician identifies the risk, the more timely the intervention can be implemented.

Phenotyping: Phenotyping can be applied in cohort construction for clinical studies, comorbidity detection and risk adjustment, quality improvement and surveillance, diagnosis and treatment planning. There are 25 different phenotypes in the dataset. We extract 41,903 patients on this task with 41.8% observed rate.

Decompensation: The objective is to determine whether a patient will decrease in the next 24 hours based on the data within a 24-hour time window. This can be utilized as an early warning in ICUs. We extract 41,744 patients. The observed rate is 43.2% and the positive rate is 3.5%.

Length of Stay: The target of length-of-stay prediction is to predict the length of stay of a patient in the ICU. This task provides a more fine-grained view of the physiological states of the patients, assisting clinicians in monitoring the progress of the disease. We obtain 33,360 patients with 43.8% observed rate. The length of stay is discretized into 10 bins, where bins 1-8 correspond to 1-8 days of stay, respectively, bin 8 for more than eight days but less than two weeks of stay, and bin 9 for living over two weeks.

A.2 Implementation Details

All experiments of this model are carried out on a Linux server equipped with RTX 2080Ti GPUs. PyTorch 2.1.2 and CUDA 12.1 deep learning libraries are applied to build and train our neural network. While training models, we apply Adam optimizer [66] with learning rate $1e-3$. We train all the models (including baselines) with the same optimizer, learning rate, and a total batch size of 256. The training epochs and hyperparameters (such as dropout and hidden dimensions) of baselines are tuned for better performance. To evaluate precisely, we repeat every experiment 3 times with different random seeds (1, 42, and 3407). The standard deviation of the experimental results appears to be slightly larger because the division of the training, validation, and test sets is related to the seed rather than being fixed. We cannot provide evaluations about imputations since we reconstruct in latent space and the absence of ground truth.

Hyperparameters: The hidden size d of all modules in SMART is 32. The heads of the multi-head attention computation in temporal attention and variable attention are set to 4. The layer number L of the MART blocks is 2. The model is pre-trained for 25 epochs and fine-tuned for 25 epochs. The unfreeze epoch is set to 5. We employ the same hyperparameter configurations across the datasets in our experiments. The probability interval p for generating masks is set to (0, 0.75), and the dropout rate is set to 0.1. The exponential moving average (EMA) decay rate is set to 0.996.

Table 3: Performance comparison of AUROC and min(Se, P+) with standard deviation on the four clinical tasks. The best results are marked in bold. The second best results are underlined.

Model	Cardiology		Sepsis	
	AUROC	min(Se, P+)	AUROC	min(Se, P+)
AdaCare	78.01±1.16	41.57±3.60	84.70±0.43	42.88±2.47
StageNet	82.55±1.75	<u>49.91±1.08</u>	91.25±1.50	56.21±3.08
ConCare	82.74±1.47	46.65±2.71	94.39±0.74	69.79±3.07
GRASP	82.98±0.66	49.53±3.23	94.43±0.76	69.75±2.10
SAFARI	81.84±1.09	47.54±2.32	<u>94.64±0.83</u>	<u>70.03±1.68</u>
PPN	83.42±1.23	48.45±1.51	<u>94.25±1.07</u>	<u>68.96±2.25</u>
RainDrop	76.86±1.67	38.31±1.41	92.90±0.79	68.74±2.42
Warpformer	83.21±2.09	48.73±3.22	93.86±0.77	65.93±1.79
PrimeNet	<u>84.19±0.52</u>	49.67±2.58	83.60±1.16	36.18±2.39
SMART	85.80±0.05	52.63±1.37	95.92±0.84	75.61±2.76

Model	In-hospital Mortality		Decompensation	
	AUROC	min(Se, P+)	AUROC	min(Se, P+)
AdaCare	78.35±1.12	42.68±0.81	87.06±1.23	49.96±1.95
StageNet	<u>85.32±0.62</u>	49.84±3.12	92.49±0.47	60.44±0.93
ConCare	84.97±0.96	49.61±1.61	92.43±0.46	61.54±1.01
GRASP	84.79±1.26	49.68±2.19	92.25±0.34	60.76±1.54
SAFARI	84.48±1.16	48.81±1.38	90.69±0.16	58.64±1.73
PPN	84.90±0.96	50.42±1.12	91.50±0.70	60.65±0.91
RainDrop	83.21±0.99	46.30±1.51	90.48±1.60	56.55±2.44
Warpformer	84.11±1.10	48.35±2.02	93.24±0.96	64.42±0.97
PrimeNet	84.82±1.05	<u>50.96±1.01</u>	91.15±1.49	54.99±2.66
SMART	85.61±0.62	51.29±1.35	94.20±0.12	65.52±1.74

Table 4: Ablation study of mask information in SMART on the Cardiology, Sepsis, and MIMIC-III in-hospital mortality.

SMART	Cardiology		Sepsis		In-hospital Mortality	
	AUPRC(%)	F1 Score(%)	AUPRC(%)	F1 Score(%)	AUPRC(%)	F1 Score(%)
Full	53.84±2.24	47.53±2.33	81.67±0.84	75.37±2.62	53.30±0.12	44.23±2.03
w/o Mask	49.03±1.93	41.98±1.71	76.35±2.76	69.51±2.11	50.43±2.44	39.70±1.60
w/o Mask in Encoder	49.74±2.43	42.70±2.11	77.37±2.53	71.96±2.26	52.29±1.06	42.65±2.91
w/o Mask in Temporal Attention	53.71±2.22	46.69±2.35	81.08±2.97	75.62±2.83	53.29±1.58	42.83±3.09
w/o Mask in Variable Attention	52.37±3.17	44.53±1.83	76.87±2.20	72.06±2.66	52.97±0.36	41.21±2.99

A.3 Experimental Results on More Metrics

In addition to AUPRC and F1 Score, we also provide the results of other metrics on the binary classification task, such as the area under the receiver operating characteristic curve (AUROC), which is a widely used metric in the healthcare field. We also provide the results of the minimum of the precision and sensitivity (min(Se, P+)), which is often used in the healthcare field [16]. The results are shown in Table 3. The results illustrate that SMART outperforms existing baselines across all metrics on the four clinical tasks. AUROC can be used to examine the model’s ability to distinguish between positive and negative samples, while min(Se, P+) can be used to evaluate the model’s ability to balance sensitivity and precision. By observing the results, we can find that the results on these metrics are consistent with the results on AUPRC and F1 Score. Specifically, AdaCare performs the worst in the baseline, and StageNet, RainDrop, and PrimeNet show unstable performance on some datasets. In summary, SMART achieves the best performance on all metrics, demonstrating its effectiveness in learning representations and predicting patient health status.

Table 5: Study on the layer numbers of the MART block on the Cardiology, Sepsis, and in-hospital mortality. The best results are marked in bold.

SMART L	Cardiology		Sepsis		In-hospital Mortality	
	AUPRC(%)	F1 Score(%)	AUPRC(%)	F1 Score(%)	AUPRC(%)	F1 Score(%)
1	52.66±2.67	44.07±2.72	81.26±2.68	74.68±2.92	53.40±1.66	42.55±2.83
2	53.84±2.24	47.53±2.33	81.67±0.84	75.37±2.62	53.30±0.12	44.23±2.03
3	52.18±1.42	46.54±1.88	81.35±3.24	74.30±2.74	53.82±1.51	43.28±3.53
4	50.56±3.90	46.65±1.75	79.98±2.91	73.68±2.02	51.89±0.71	44.97±2.62

Table 6: Study on the mask ratios of SMART in the pre-training stage on the Cardiology, Sepsis, and in-hospital mortality. The best results are marked in bold.

SMART p	Cardiology		Sepsis		In-hospital Mortality	
	AUPRC(%)	F1 Score(%)	AUPRC(%)	F1 Score(%)	AUPRC(%)	F1 Score(%)
(0, 0.25)	52.81±2.63	44.05±2.19	80.78±2.46	74.46±2.95	52.55±0.39	44.17±3.91
(0, 0.5)	53.77±2.45	46.92±0.61	80.89±2.23	74.04±2.75	53.41±0.78	44.90±3.56
(0, 0.75)	53.84±2.24	47.53±2.33	81.67±0.84	75.37±2.62	53.30±0.12	44.23±2.03
(0.25, 0.75)	53.66±2.64	45.06±3.09	81.82±2.04	74.71±3.03	53.48±0.66	43.03±3.00
(0.5, 0.75)	53.63±2.38	46.00±1.64	81.69±2.02	74.87±2.29	53.37±1.00	42.45±2.39

A.4 Ablation Study on the Mask Information

We conduct an ablation study on the mask information in SMART. This ablation study can be regarded as an extension of the reduced version **w/o Mask** in Table 2 and it is a deeper exploration of how the mask information affects the performance of SMART. Specifically, we remove the mask information in the input encoder and the attentions in the MART blocks separately, namely **w/o Mask in Encoder**, **w/o Mask in Temporal Attention**, and **w/o Mask in Variable Attention**. The results are shown in Table 4. We observe that the performance of SMART drops significantly when the mask information is removed in the input encoder, which indicates that the mask information in the input encoder is critical to learn complete representations. When the mask information is removed in the attentions, the performance of SMART also decreases, which demonstrates that the mask information in the attentions is essential to resist noise when capturing the temporal and variable interactions. These results of the ablation study on the mask information further confirm each design in SMART is indispensable and contributes to the improvement.

A.5 Study on the Hyperparameters

To explore the affect of the hyperparameters on the performance of SMART, we conduct a study on the layer numbers of the MART block and the mask ratios in the pre-training stage. The results are shown in Table 5 and Table 6. On the layer numbers of the MART block, the results show that the performance of SMART is the best when the layer numbers are set to 2. Besides, we find that too much layer numbers may lead to decreased performance.

When it comes to the mask ratios in the pre-training stage, when the mask ratios are set to (0, 0.75), SMART achieves the best performance on the Cardiology and Sepsis, whereas the mask ratio $p=(0.25, 0.75)$ obtains the best AUPRC and $p=(0, 0.5)$ achieves the best F1 Score on the in-hospital

Table 7: Performance comparison with standard deviation on Cardiology, Sepsis, and in-hospital mortality with basic baselines.

Model	Cardiology		Sepsis		In-hospital Mortality	
	AUPRC(%)	F1 Score(%)	AUPRC(%)	F1 Score(%)	AUPRC(%)	F1 Score(%)
GRU	47.10±3.56	37.84±0.92	49.73±3.12	47.66±2.15	50.68±2.07	41.70±2.21
Transformer	42.92±3.70	42.82±3.73	22.56±4.41	24.75±3.90	43.07±3.68	42.94±3.36
SMART	53.84±2.24	47.53±2.33	81.67±0.84	75.37±2.62	53.30±0.12	44.23±2.03

mortality. These experiments shed light on the importance of hyperparameters in the performance of SMART, and provide guidance for the selection of hyperparameters in practice.

A.6 Comparisons with Basic Baselines

We compare SMART with basic baselines on the Cardiology, Sepsis, and in-hospital mortality tasks, including gated recurrent units (GRU) and Transformer [22]. The results are shown in Table 7. We observe that SMART outperforms the basic baselines on all datasets, which shows the effectiveness of SMART in learning representations and predicting patient health status. The results also show that the basic baselines have poor performance especially on the Sepsis datasets, which may be due to the complexity and higher missing rate of the dataset.

B Broader Impact

The proposed model is designed to improve the representation learning capability of EHR data in ICU scenarios. The model can be used to predict patients' health status, such as in-hospital mortality, sepsis, and other diseases. The model can be used to assist physicians in making clinical decisions, such as early warning of patients' health status, which may help reduce the mortality rate of patients. The model can also be used to assist in the allocation of medical resources, such as the allocation of ICU beds, which may help reduce the burden on the healthcare system. However, the model may also have some negative impacts. For example, the model's predictions may be biased, which may result in patients being treated unfairly. Additionally, there may be ethical issues with this model, as it may lead to decisions that are not in the best interest of the patient. Therefore, it is important to carefully evaluate the model before actually using it. It should be noted that the model cannot replace doctors in making decisions. The model is only a tool to assist doctors in making decisions. The model should be used in conjunction with doctors to make the best decisions for patients.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We focus on analyzing EHR data from ICU and propose a novel model using missing-aware and self-supervised pre-training to enhance representation learning capability.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discuss the limitations of the proposed model in Section 5.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: This paper does not include theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provide detailed information on hyperparameters in Appendix A.2. Besides, we provide the code for the experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide the code link in the paper.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We provide detailed information on hyperparameters in Appendix A.2.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We repeat each experiment with three random seeds and report both the mean and standard deviation of the results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute in Appendix A.2.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: The research conducted in the paper conforms with the NeurIPS Code of Ethics. We conducted our study using publicly available EHR data, which had passed ethical review.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We discuss the potential positive societal impacts and negative societal impacts of the proposed model in Appendix B.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper does not release data or models that have a high risk for misuse.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We cite the original papers that produced the datasets.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: We release our code and provide instructions on how to reproduce the results.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.