



SepsisLab: Early Sepsis Prediction with Uncertainty Quantification and Active Sensing

Changchang Yin
yin.731@osu.edu
The Ohio State University
Columbus, Ohio, USA

Pin-Yu Chen
pin-yu.chen@ibm.com
IBM Research
Yorktown Heights, New York, USA

Bingsheng Yao
b.yao@northeastern.edu
Northeastern University
Boston, Massachusetts, USA

Dakuo Wang
d.wang@northeastern.edu
Northeastern University
Boston, Massachusetts, USA

Jeffrey Caterino
jeffrey.caterino@osumc.edu
The Ohio State University Wexner
Medical Center
Columbus, Ohio, USA

Ping Zhang*
zhang.10631@osu.edu
The Ohio State University
Columbus, Ohio, USA

ABSTRACT

Sepsis is the leading cause of in-hospital mortality in the USA. Early sepsis onset prediction and diagnosis could significantly improve the survival of sepsis patients. Existing predictive models are usually trained on high-quality data with few missing information, while missing values widely exist in real-world clinical scenarios (especially in the first hours of admissions to the hospital), which causes a significant decrease in accuracy and an increase in uncertainty for the predictive models. The common method to handle missing values is imputation, which replaces the unavailable variables with estimates from the observed data. The uncertainty of imputation results can be propagated to the sepsis prediction outputs, which have not been studied in existing works on either sepsis prediction or uncertainty quantification. In this study, we first define such propagated uncertainty as the variance of prediction output and then introduce uncertainty propagation methods to quantify the propagated uncertainty. Moreover, for the potential high-risk patients with low confidence due to limited observations, we propose a robust active sensing algorithm to increase confidence by actively recommending clinicians to observe the most informative variables. We validate the proposed models in both publicly available data (i.e., MIMIC-III and AmsterdamUMCdb) and proprietary data in The Ohio State University Wexner Medical Center (OSUWMC). The experimental results show that the propagated uncertainty is dominant at the beginning of admissions to hospitals and the proposed algorithm outperforms state-of-the-art active sensing methods. Finally, we implement a SepsisLab system for early sepsis prediction and active sensing based on our pre-trained models. Clinicians and potential sepsis patients can benefit from the system in early prediction and diagnosis of sepsis.

*Corresponding Author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '24, August 25–29, 2024, Barcelona, Spain

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0490-1/24/08

<https://doi.org/10.1145/3637528.3671586>

CCS CONCEPTS

• Information systems → Data mining; • Applied computing → Health informatics.

KEYWORDS

Clinical decision support, Early sepsis prediction, Active sensing, Electronic health record, Deep learning

ACM Reference Format:

Changchang Yin, Pin-Yu Chen, Bingsheng Yao, Dakuo Wang, Jeffrey Caterino, and Ping Zhang. 2024. SepsisLab: Early Sepsis Prediction with Uncertainty Quantification and Active Sensing. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24)*, August 25–29, 2024, Barcelona, Spain. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3637528.3671586>

1 INTRODUCTION

Sepsis, defined as life-threatening organ dysfunction in response to infection, contributes to up to half of all hospital deaths and is associated with more than \$24 billion in annual costs in the United States [13]. Existing studies [14] have shown that a sepsis patient may benefit from a 4% higher chance of survival if they are diagnosed 1 hour earlier, so developing an early sepsis onset prediction system can significantly improve clinical outcomes.

Existing machine-learning-based predictive models [7, 10, 21, 38] are usually trained on high-quality data with few missing information, while missing values widely exist in emergency department (ED) and emergency medical services (EMS) settings, which would cause most existing sepsis prediction models to suffer from performance decline and high uncertainty. In addition, existing studies [22, 30] have shown that for sepsis cases, most patients have already progressed into sepsis before the admissions to hospitals or during the first hours of admissions. Thus it is critical to develop accurate sepsis prediction systems that can handle high missing-rate settings (e.g., cold-start setting with only several limited vital signs).

A common method to handle missing variables is imputation, in which missing values are replaced by estimates from the observed data. To use the existing methods, we will need data imputations, which come with a new problem for the downstream sepsis prediction tasks: the uncertainty of imputation results can propagate to the sepsis prediction models. Especially for deep learning models, a small perturbation in the input variables might cause a significant




Figure 1: Workflow of SepsisLab system.

change in the predicted risk [15, 20]. When the prediction models are sensitive to the highly uncertain input (i.e., imputed variable), the generated outputs are not reliable, so it is critical to quantify and reduce such kind of uncertainty. However, unlike epistemic uncertainty [28] and aleatory uncertainty [4], the propagated uncertainty from the (imputed) input has not been investigated.

In this study, we develop an early sepsis prediction system SepsisLab that can quantify and reduce such kind of propagated uncertainty from missing value and imputation. Figure 1 displays the workflow of SepsisLab system. Given a patient’s data with limited observations, we first adopt an imputation model to estimate the distribution (i.e., mean and standard deviation) of missing values. The standard deviation can be treated as the uncertainty of the imputed results. Then we propose a time-aware sepsis prediction model to predict whether the patients will suffer from sepsis in the coming hours. The prediction model can generate sepsis risk and uncertainty simultaneously. Given the estimated uncertainty, we further propose a robust active sensing algorithm to recommend clinicians observe the most informative lab test items that can maximally reduce the uncertainty for the potential high-risk patients. The active sensing module can significantly improve downstream sepsis prediction performance by providing more accurate observation and reducing the propagated uncertainty.

To demonstrate the effectiveness of the proposed models, we conduct experiments on real-world clinical datasets (including two publicly available datasets MIMIC-III [9] and AmsterdamUM-Cdb [29], and proprietary data from The Ohio State University Wexner Medical Center (OSUWMC)). Experimental results show that the developed system can successfully work on both high- and low-missing-rate settings and achieve state-of-the-art sepsis prediction performance. Finally, we develop a SepsisLab system for deployment to integrate into clinicians’ workflow, which paves the way for human-AI collaboration and early intervention for sepsis management.

We summarize our contributions as follows:

- We introduce propagated uncertainty to deep learning models, a new source of uncertainty different from widely studied aleatoric uncertainty and epistemic uncertainty
- We adopt uncertainty propagation to successfully qualify the propagated uncertainty, and the experimental results demonstrate the propagated uncertainty is dominant at the beginning of patients’ admissions to hospital.
- We propose a new active sensing framework RAS, which could effectively select variables to observe, and the experiments demonstrate the effectiveness of the proposed propagated uncertainty qualification method.
- We design an interactive system SepsisLab¹ to make clinicians able to easily use and effectively interact with the models.

2 RELATED WORK

In this section, we briefly review the existing studies related to sepsis prediction systems, uncertainty qualification and active sensing.

2.1 Sepsis Prediction Systems

Sepsis is a heterogeneous clinical syndrome that is the leading cause of mortality in hospital intensive care units (ICUs) [24, 33]. Early prediction and diagnosis may allow for timely treatment and lead to more targeted clinical interventions. Screening tools have been used clinically to recognize sepsis, including qSOFA [25], MEWS [27], NEWS [26], and SIRS [3]. However, those tools were designed to screen existing symptoms as opposed to explicitly early predicting sepsis before its onset, and their efficacy in sepsis diagnosis is limited. With recent advances, deep learning methods have shown great potential for accurate sepsis prediction [7, 10, 21, 38]. Although the methods achieved superior performance, they face a critical limitation: the models need to take the complete observation of a list of variables (including vital signs and lab tests), while lots of variables are missing in real-world data (especially in the first hours of admissions). Existing studies [7, 10, 38] usually impute the missing values before the prediction, which raises a new problem that the sepsis prediction models will heavily rely on the imputation methods. The imputation uncertainty would also be propagated to downstream prediction models. Thus it is necessary to quantify the propagated uncertainty, especially for high-stakes sepsis prediction tasks.

2.2 Uncertainty Qualification

Understanding what a model does not know is a critical part of many machine learning systems. Despite the superior performance deep learning models have achieved in various domain, they are usually over-confident about the predictions, which could limit their applications to real-world risk-sensitive settings (e.g., in health-care). Uncertainty quantification methods play a pivotal role in reducing the impact of uncertainties during both optimization and decision making processes [1]. Existing uncertainty qualification work [4, 6, 11, 23] has widely studied epistemic uncertainty and

¹<https://github.com/yinchangchang/SepsisLab>




Figure 2: Model framework. (A) The imputation model takes observed variables and corresponding timestamps as input, and generates the distribution of missing values. (B) Sepsis prediction model produces the patients’ sepsis onset risks with uncertainty based on the imputed data. (C) shows the uncertainty quantification method with Monte-Carlo sampling. (D) displays the uncertainty propagation method that can estimate propagated uncertainty by multiplying models’ gradient over imputed variables and the imputation uncertainty.

aleatoric uncertainty. However, most existing uncertainty qualification studies ignore an important uncertainty source: the uncertainty propagated from the uncertainty of input (e.g., widely existing missing values). In this study, we aim to investigate and reduce the propagated uncertainty.

2.3 Active Sensing

Active sensing aims to improve the target tasks’ performance by actively selecting most informative variables with the minimal cost. Yu et. al [37] propose to select the informative variables based on mutual information and predictive variance. However, the model is based on Bayesian co-training framework, the prediction ability of which is not as good as deep neural networks when handling large-scale time serial data. Yoon et. al [36] attempt to solve the active sensing problem by proposing an RNN-based model (i.e., Deep Sensing). The Deep Sensing framework involves learning 3 different networks: an interpolation network, a prediction network and an error estimation network. Each network is separately optimized for its own objective and then combined together after training to be used for active sensing. Jarrett et. al [8] propose an Inverse Active Sensing (IAS) to require negotiating (subjective) trade-off between accuracy, speediness, and cost of information. Yoon et. al [35] propose an RL-based framework (Active Sensing using Actor-Critic models, ASAC) to directly optimize the predictive power after active sensing. Although the methods achieved superior performance in the target prediction tasks, they failed to measure the uncertainty of both missing values and model output risks, which limit their application in high-stakes clinical settings.

In this study, we aim to develop an accurate sepsis prediction system with propagated uncertainty quantification and incorporate active sensing algorithms to reduce the propagated uncertainty.

3 METHODOLOGY

In this section, we present the proposed sepsis prediction system SepsisLab, including a missing value imputation model, an early sepsis prediction model, and an active sensing algorithm.

3.1 Notation and Problem Statement

In this study, we aim to predict sepsis onset with limited clinical variables observed. We consider the following setup. A patient has a sequence of clinical variables (i.e., lab test data and vital sign data) with timestamps. Let $Z \in \{R \cup *\}^{n \times k}$ denote the observations of variables, where $*$ represents missing values, n denotes the number of collections of observations and k denotes the number of unique clinical variables. $T \in R^n$ denotes the observation timestamps. $Y \in \{0, 1\}^n$ denotes the ground truth of whether the patient will progress to sepsis in the coming hours. Following [10, 38], we set the prediction window as 4 hours. Due to the existence of missing values, we impute the missing values first and use $X \in R^{n \times k}$ to denote the imputed results.

Given a loss function $\mathcal{L}$ and a distribution over pairs (X, Y) , the goal is to find a function f that minimize the expected loss:

$$f^* = \arg \min_f E[\mathcal{L}(f(X), Y)] \quad (1)$$

We list the important notations in Table 1.

Table 1: Basic Notations.

Notation	Description
Z	Observed variables with missing values.
X	Imputed variables.
Y	Labels for sepsis prediction.
p_i	Predicted sepsis risk at i^{th} collection.
T	Timestamps for observations.
M	Masking indicator for imputation training.
n	The number of collections of variables.
k	The number of unique variables.
e_i^t	Time embedding vector for i^{th} collection.
e_i	The embedding for i^{th} collection.
μ	Mean of missing values.
σ	Standard deviation of missing values.
w_*, b_*	Learnable parameters.
s_i	Hidden state of imputation model.
h_i	Hidden state of the sepsis prediction model.
U	Computed uncertainty.
U_x	Propagated uncertainty.
U_w	Epistemic uncertainty.
ρ_{ij}	Correlation between i^{th} and j^{th} variable.
lr	Learning rate.

3.2 Missing Value Imputation

We assume the missing values follow the Gaussian distributions and impute the missing values by estimating the distribution of variables (i.e., the mean and covariance). Figure 2(A) shows the framework of our imputation model.

Following [31], we first use mean-imputation to preprocess the observational data Z and send the embedding of Z to LSTM to model the patient's health states.

Embedding layer. In the i^{th} collection, we have observational values Z_i , observation time T_i . We use a fully connected layer to embed the observed variable in the collection:

$$e_i = w_e[Z_i; e_i^t] + b_e, \quad (2)$$

where $[\bullet; \bullet]$ denotes concatenation operation. $w_e \in R^{(k+2d) \times d}$ and $b_e \in R^d$ are learnable variables. $e_i^t \in R^{2d}$ denotes the time embedding and is computed as follows:

$$e_{i,j}^t = \sin\left(\frac{T_i * j}{T_{max} * d}\right), e_{i,d+j}^t = \cos\left(\frac{T_i * j}{T_{max} * d}\right), \quad (3)$$

where $0 \leq j < d$, and T_{max} denotes the max value of T .

Time-aware LSTM encoder. Given the embedding vectors $[e_1, e_2, \dots, e_n]$, we use LSTM to model the patients' states:

$$s_1, s_2, \dots, s_T = LSTM(e_1, e_2, \dots, e_n) \quad (4)$$

Missing value distribution estimation. A fully connected layers is used to generate the parameters of the missing value distribution:

$$\mu_i = w_\mu s_i + b_\mu, \quad \sigma_i = ReLU(w_\sigma s_i + b_\sigma), \quad (5)$$

where $w_\mu, w_\sigma \in R^k$ and $b_\mu, b_\sigma \in R$ are learnable variables.

We train the imputation model with the mean square error loss function:

$$\mathcal{L}_{imp}(Z, M, \mu) = \sum_{i=1}^n \sum_{j=1}^k M_{i,j} (\mu_{i,j} - Z_{i,j})^2, \quad (6)$$

where $M \in \{0, 1\}^{n \times k}$ denotes the indices of masked variables. $M_{i,j}$ is 1 if the j^{th} variable in i^{th} collection is observed and masked; otherwise, 0. Replacing the missed values $*$ with the estimates μ , the observed variables Z become $X \in R^{T \times n}$.

After the imputation model is well-trained with Equation 6, we further learn to estimate the standard deviation σ by finetuning w_σ and b_σ and fixing other parameters. We minimize the following loglikelihood loss:

$$\mathcal{L}_\sigma(Z, M, \mu, \sigma) = \sum_{i=1}^n \sum_{j=1}^k M_{i,j} \left[\frac{(\mu_{i,j} - Z_{i,j})^2}{2\sigma_{i,j}^2} + \frac{\sigma_{i,j}^2}{2} \right] \quad (7)$$

3.3 Sepsis Prediction Model

With the imputed results to replace the missing values, we continue to predict whether the patients will suffer from sepsis in the coming hours. The framework of sepsis prediction model is shown in Figure 2(B).

Similar to Equation 8 in the imputation model, we use the same embedding layers in the imputation model.

$$e_i' = w_e[X_i; e_i^t] + b_e, \quad (8)$$

where the time embedding e_i^t is the same as in Equation 3.

Then we use LSTM [5] to model the patient's health states. A fully connected layer and a Sigmoid layer is followed to generate the sepsis risks:

$$p_i = \text{Sigmoid}(w_s h_i + b_s), \text{ where } t = 1, 2, \dots, T \quad (9)$$

$$h_1, h_2, \dots, h_n = LSTM(e_1', e_2', \dots, e_n'), \quad (10)$$

where $w_s \in R^d$ and $b_s \in R$ are learnable parameters.

The model is trained by minimizing the binary cross-entropy loss:

$$\mathcal{L}_{cls}(p, Y) = \frac{1}{n} \sum_{i=1}^n -y_i \log(p_i) - (1 - y_i) \log(1 - p_i) \quad (11)$$

3.4 Sources of Uncertainty

When applying deep learning methods to high-stakes sepsis prediction tasks, the lack of uncertainty quantification will make the models less reliable. In this subsection, we investigate two main sources of uncertainty.

Uncertainty from the model parameters. Existing uncertainty qualification work [4, 6, 11, 23] has widely studied epistemic uncertainty, which accounts for uncertainty in the model parameters, especially for the huge amount of parameters in deep learning models. Following [11], we use drop-out during the test phase and run the inference many times to quantify such kind of uncertainty.

Uncertainty from missing values. Superior risk prediction models in the healthcare domain heavily rely on high-quality complete input. However, missing values (e.g., vital signs and lab test results) widely exist in real-world clinical settings. Most risk prediction methods [7, 10, 21, 38] first impute the missing values and then make predictions based on the imputed values. The accuracy

of the imputation methods can directly affect the performance of the predicted sepsis risks. The uncertainty from the imputation results can be directly propagated to downstream sepsis prediction models.

3.5 Uncertainty Definition

We use the variance of prediction models' output to define the two kinds of uncertainty mentioned above. Patients' data X contains a sequence of collections of variables. We can use all the observations until the current collections to make predictions. When applying active sensing algorithms to reduce the propagated uncertainty with additional observations, we can only request the variables in the current collection.

In the active sensing task, we only focus on the uncertainty related to the latest collection. In the following subsections, for simplicity, at a given time T_i , we use x to represent the i^{th} observation (i.e., X_i), and use $f_w(x)$ rather than $f_w(X)$ to denote the predicted risk, where w means all the learnable parameters in the sepsis prediction model.

We assume the input variables $x \in R^k$ and model parameters w follow Gaussian distributions $\mathcal{N}(\mu_x, \sigma_x)$ and $\mathcal{N}(\mu_w, \sigma_w)$. $\mu_x \in R^k$ and $\sigma_x \in R^k$ can be estimated with Equation 5. Let y denote the sepsis prediction label for the patient at current time.

Following existing studies [11], we define the uncertainty of predicted risk as the variance of model outcomes:

$$U = \int_w \int_x (f_w(x) - \mu_y)^2 \rho(x) dx \rho(w) dw = U_x + U_w \quad (12)$$

$$\text{where } U_x = \int_w \int_x (f_w(x) - \mu_{y_w})^2 \rho(x) dx \rho(w) dw,$$

$$U_w = \int_w (\mu_{y_w} - \mu_y)^2 \rho(w) dw,$$

$$\mu_{y_w} = \int_x f_w(x) \rho(x) dx,$$

$$\mu_y = \int_w \int_x f_w(x) \rho(x) \rho(w) dx dw,$$

where $\rho(\bullet)$ denotes the density function.

We split the uncertainty into two terms. The second term U_w is caused by the model uncertainty from the model parameters, so we just focus on the first term U_x when actively selecting unobserved variables.

When the model parameter w is fixed, we can estimate the propagated uncertainty as:

$$U_x^{(w)} = \int_x (f_w(x) - \mu_{y_w})^2 \rho(x) dx \quad (13)$$

3.6 Propagated Uncertainty Quantification

3.6.1 Propagated Uncertainty for Linear Target Prediction. When the sepsis risk prediction function is a linear function, $f_w(x) = \sum_j w_j x_j$, following [12], we compute the uncertainty in Equation 13 as:

$$U_x^{(w)} = \sum_i w_i^2 \sigma_{x_i}^2 + \sum_i \sum_{j \neq i} w_i w_j \rho_{ij} \sigma_{x_i} \sigma_{x_j}, \quad (14)$$

where ρ_{ij} denotes the correlation between i^{th} and j^{th} variable. It is easy to compute the propagated uncertainty for linear function based on Equation 14 for linear function. The calculation details

Algorithm 1 Adversarial Training

Input: observations X , missing value distribution μ_x, σ_x , outcome Y , step size s_{adv} , step n_{adv} , learning rate lr ;

```

1: repeat
2:   Sample a batch of patients' data,  $x, \sigma_x, y$ ;
3:   Initialize the perturbation  $\delta$  with Gaussian distribution  $N(0, \sigma_x)$  and constraint  $-2\sigma_x < \delta < 2\sigma_x$ ;
4:   Compute  $f_w(x)$  and the first order gradient  $\nabla_x$ ;
5:   for  $i = 1, \dots, n_{adv}$  do
6:     Calculate  $g = \nabla_\delta g(\delta, x)$ 
7:     Update  $\delta = \delta + s_{adv} \times g$ 
8:   end for
9:   Calculate loss  $L$  in Equation 19 and gradient  $\nabla_w L$ ;
10:  Update  $w = w - lr \times \nabla_w L$ ;
11: until Convergence.
```

for Equation 14 can be found in subsection A.1 in supplementary materials.

The propagated uncertainty reduction after observing i^{th} variable is:

$$U_x^{(w)}(i) = w_i^2 \sigma_{x_i}^2 + \sum_{j \neq i} w_i w_j \rho_{ij} \sigma_{x_i} \quad (15)$$

3.6.2 Propagated Uncertainty for Non-Linear Target Prediction. For the non-linear sepsis prediction function, we use the Taylor expansion as approximate function:

$$\tilde{f}_w(x + \delta) = f_w(x) + \delta^T \nabla_x f_w(x) \quad (16)$$

We can use the uncertain propagation in Equation 14 as the approximation of the uncertainty of non-linear function f_w . However, the propagated uncertainty estimation for non-linear functions are biased on account of using a truncated series expansion. The extent of this bias depends on the nature of the function.

The absolute difference between the two values $\tilde{f}_w(x + \delta)$ and $f_w(x + \delta)$ is:

$$g(\delta, x) = |f_w(x + \delta) - f_w(x) + \delta^T \nabla_x f_w(x)| \quad (17)$$

When $g(\delta, x)$ is small enough in the neighborhood near μ_x (i.e., f_w is locally linear), the propagated uncertainty in Equation 14 is still accurate and able to guide the active sensing.

3.7 Robust Active Sensing

3.7.1 Adversarial Training for Local Linearity. Existing studies [18, 19] have shown that adversarial training can encourage the local linearity of the learned functions. In this study, we adopt adversarial training to make the target prediction function locally linear in a neighborhood near the mean value of input x .

$$\mathcal{L}_{adv} = \min_w \max_\delta g(\delta, \mu_x), \text{ where } -2\sigma_x < \delta < 2\sigma_x \quad (18)$$

The risk prediction model is trained with a weighted sum of classification loss and adversarial loss:

$$\mathcal{L} = \alpha \mathcal{L}_{cls} + (1 - \alpha) \mathcal{L}_{adv}, \quad (19)$$

where $0 < \alpha < 1$ is a hyper-parameter.

We consider the quantity:

$$\gamma(\sigma, x) = \max_{-2\sigma \leq \delta \leq 2\sigma} |f_w(x + \delta) - f_w(x) - \delta^T \nabla_x f_w(x)| \quad (20)$$

to be a measure of how linear the surface is within a neighborhood near x . We call this quantity the local linearity measure. The missing variables follow Gaussian distribution, so δ lies within two standard deviations with more than probability 95%. The uncertainty estimation error would be less than $\gamma(\sigma_x, x)$ with probability more than 95%.

Algorithm 1 describes the training process of the sepsis prediction model.

3.7.2 Active Sensing. The approximate uncertainty of the risk prediction outcome is defined as:

$$U_x^{(w)} = \sum_i \nabla_{x_i}^2 \sigma_{x_i}^2 + \sum_i \sum_{j \neq i} \nabla_{x_i} \nabla_{x_j} \rho_{ij} \sigma_{x_i} \sigma_{x_j} \quad (21)$$

The propagated uncertainty reduction after observing i^{th} variable is:

$$U_i^{(w)} = \nabla_{x_i}^2 \sigma_{x_i}^2 + \sum_{j \neq i} \nabla_{x_i} \nabla_{x_j} \rho_{ij} \sigma_{x_i} \sigma_{x_j} \quad (22)$$

Considering the distribution of w , we use Monte-Carlo dropout to sample model parameters and use the average uncertainty of $U_x^{(w)}(i)$ to approximately compute $U_x(i)$:

$$U_x(i) = \int_w U_x^{(w)}(i) \rho(w) dw \quad (23)$$

We can select the unobserved variables based on the maximal uncertainty criterion.

$$i^* = \arg \max_i U_x(i), \quad (24)$$

where i^* is the best variable to observe. Figure 2(D) shows the workflow of propagated uncertainty quantification and active sensing methods.

4 EXPERIMENT SETUP

To demonstrate the effectiveness of the proposed method, we conducted experiments on real-world datasets.

4.1 Datasets

Datasets. We validate our system on two publicly available datasets (MIMIC-III² and AmsterdamUMCdb³) and one proprietary dataset extracted from OSUWMC⁴. We first extracted all the sepsis patients with sepsis-3 criteria [24] in the datasets. For each sepsis patient, we select 1 control patient with the same demographics (i.e., age and gender). We extracted 26 vital signs and lab tests from the datasets. A detailed list of clinical variables can be found in supplementary materials. The statistics of the three datasets are displayed in Table 2.

Variables Used for Sepsis Prediction. Following [33], we use following variables to model sepsis patients' health states: heart rate, Respiratory, Temperature, Spo2, SysBP, DiasBP, MeanBP, Glucose, Bicarbonate, WBC, Bands, C-Reactive, BUN, GCS, Urineoutput,

Table 2: Statistics of MIMIC-III and AmsterdamUMCdb

	MIMIC	AmsterdamUMCdb	OSUWMC
#. of patients	21,686	6,560	85,181
#. of male	11,862	3,412	41,710
#. of female	9,824	3,148	43,471
Age (mean $\pm$ std)	60.7 $\pm$ 11.6	62.1 $\pm$ 12.3	59.3 $\pm$ 16.1
Missing rate	65%	68%	75%
Sepsis rate	32%	35%	29%




Figure 3: Settings of sepsis onset prediction.

Creatinine, Platelet, Sodium, Hemoglobin, Chloride, Lactate, INR, PTT, Magnesium, Aniongap, Hematocrit, PT.

The first 8 variables are immediately available before vital signs. The missing rates of the variables can be found in Table 7 in subsection A.2.

4.2 Setup

We mimic the cold-start environment where only vital signs are immediately available, while all the lab tests can be observed after the assignment. Figure 3 displays the setting of the experiments. After the patients arrive at the hospital, we start to predict whether the patients will suffer from sepsis in 4 hours. We run the prediction process hourly until the patients have been diagnosed with sepsis or discharged. When the model's output has a high uncertainty due to the limited observations, the active sensing algorithms can select the missing lab tests to observe. Based on the lab testing turnaround times policy of OSUWMC, most lab results will be available in less than 30~60 min⁵ (or even sooner for sepsis patients with high priority), so the observation results for the selected lab items can be used in the same hour to update the predicted sepsis risk. Note that when active sensing algorithms select some variables that are not collected at the corresponding time, we use the estimates from other observed variables as the active observation results.

4.3 Methods for Comparison

We compare the proposed model with following methods:

- **Random sensing:** We randomly select the masked values to observe for random sensing.
- **Active sensing MI [37]:** The method selects the most informative variables based on the mutual information.

²<https://mimic.physionet.org/>

³<https://amsterdammedicaldatascience.nl>

⁴<https://wexnermedical.osu.edu/>

⁵ <https://rb.gy/s4jiif>

- **Virtual adversarial training (VAT)** [17]: VAT proposes to make the learned function locally linear with local a smoothness regularization method. Then we use the same variable selection method as ours to select missing values.
- **Monte Carlo sampling**: Existing studies [11] use Monte-Carlo dropout to measure the epistemic uncertainty. Similarly, we use Monte-Carlo sampling to estimate the propagated uncertainty by sampling the values of the unobserved variables based on the Gaussian distribution and select the variable with maximal variance in generated output, as Figure 2(C) shows.
- **Robust active sensing (RAS)**: RAS is the proposed method. To demonstrate the effectiveness of the adversarial training, we implement three versions of RAS. RAS is the main version. RAS^L uses the linear constraint to make the learned function locally linear. RAS^N means the model is only trained by minimizing the classification loss in Equation 11 without any linearity constraint.

For fair comparison to the baselines, all active sensing algorithms use the same deep-learning sepsis prediction model backbone. Our previous works [33, 38] have shown that LSTM can successfully model the time series EHR data and achieve superior performance in the sepsis prediction tasks, so we use LSTM as the model backbone. Note that the proposed active sensing methods are generalizable to various deep learning frameworks.

5 RESULTS

We now report the performance of SepsisLab in the three datasets. We focus on answering the following research questions by our experimental results:

- **Q1: How does the model uncertainty affect the sepsis prediction performance?**
- **Q2: How does the active sensing algorithm reduce the propagated uncertainty?**
- **Q3: How does the active sensing algorithm improve the sepsis prediction performance?**

5.1 Q1: How does the model uncertainty affect the sepsis prediction performance?

The existence of uncertainty makes AI models less reliable and less accurate when applying the models to real-world high-stakes scenarios. In this subsection, we aim to show how the model uncertainty affects sepsis prediction performance by analyzing the relation between uncertainty and prediction performance.

5.1.1 Prediction Performance over Uncertainty Scales. We compute the uncertainty of the sepsis onset prediction model's output with Equation 12 and split the patients into 6 sets with different uncertainty scales. Then we calculate the sepsis onset prediction performance on AUROC inside each set. Figure 4 displays the model performance over the different uncertainty scales in the three datasets. We conducted experiments in two settings. In active sensing setting, we compute the AUROC after active sensing algorithms are used. In the observed data setting, we directly run the data in the observed data (including all the recorded vital signs and lab tests) and compute the AUROC. The results in both settings show that when uncertainty is higher, the model performance becomes less




Figure 4: Sepsis onset prediction performance with different uncertainty.




Figure 5: Uncertainty distribution over times after admission.




Figure 6: Uncertainty over different active sensing ratios.

accurate, so a good active sensing framework can improve the prediction performance by reducing the uncertainty of the prediction model's output.

5.1.2 Uncertainty Scales over Time. We quantify the model uncertainty at different times from admissions. Figure 5 displays the average uncertainty scales.

Figure 5 shows that in the first 15 hours, propagated uncertainty is dominant in sepsis onset risk prediction models. We speculate the reason is that at the beginning most variables have not been observed and the missing values cause the main uncertainty, which is consistent with our clinical experts' experience. With more variables collected, the propagated uncertainty decreases a lot after 15 hours of the admissions.

Because the missing variables can cause high uncertainty during the first hours, it is critical to quantify the propagated uncertainty when applying risk prediction models to high missing-rate settings.

5.2 Q2: How does the active sensing algorithm reduce the propagated uncertainty?

Based on the estimated uncertainty, we propose active sensing algorithms to further reduce the prediction uncertainty by recommending clinicians collect more unobserved variables. We conduct experiments to show whether uncertainty can be significantly reduced with minimal additional variables observed.




Figure 7: Inference time cost over times after admission.

5.2.1 Uncertainty with Different Active Sensing Ratio. Figure 6 displays the average uncertainties for sepsis prediction results with different active sensing ratios. The results show that with more missing variables observed, the uncertainty on the predicted sepsis risks are significantly reduced. Besides, all the versions of the proposed RAS reduce more uncertainty than the baselines, which demonstrates the effectiveness of the proposed active sensing algorithms on uncertainty reduction.

5.2.2 Uncertainty Quantification Efficiency. We also investigate the time cost for uncertainty quantification during the inference phase. Figure 7 displays the inference time cost for uncertainty quantification. The results show that RAS can achieve much less time than the baselines, which makes the SepsisLab system work more efficiently during the active sensing phase.

5.3 Q3: How does the active sensing algorithm improve the sepsis prediction performance?

The goal of SepsisLab is to accurately predict the sepsis so as to provide reliable decision-making support to clinicians. We conduct experiments to show sepsis prediction performance improvement with the active sensing algorithms.

5.3.1 Sepsis onset Prediction Results. Table 3 displays the risk prediction performance with different active sensing ratios (i.e., 2%-8%). With additional variables observed, all the methods can achieve more accurate prediction performance for sepsis onset. Moreover, all the active sensing algorithms outperform the random sensing baseline with the same observation rate, which demonstrates that active sensing can improve downstream tasks' performance. Among the active sensing algorithms, the proposed RAS achieved the best performance with different active sensing ratios, which demonstrate the effectiveness of the proposed model.

5.3.2 Ablation Study. We have three versions of the framework. RAS^N directly uses the gradient to estimate propagated uncertainty. RAS^L uses a linear regularization term to make the model locally smooth, while RAS uses adversarial training. For RAS^L and RAS versions, the additional terms change the loss functions. We conduct experiments to show whether the additional terms can improve model training. We train the three versions of models independently and test them on all the observed data (without active sensing). As Table 4 shows, RAS^L and RAS outperform RAS^N , which demonstrates local linearity can further improve prediction performance. With adversarial training, RAS can achieve better local linearity than RAS^L and thus perform the best, which also explains why the RAS outperforms better than Monte-Carlo sampling in Table 3 in the active sensing.




Figure 8: User Interface of Our SepsisLab System. (A) Patient list with sepsis risk prediction score. (B) The patient's demographics and the dashboard of the patient's historical observations. (C) Predicted sepsis risk score with uncertainty range and recommended lab test items to observe.

We also conduct more experiments with different backbones (e.g., RNN, GRU, FC) and display the performance in Table 6 in subsection A.3. The experimental results show that the proposed model can consistently improve the prediction performance for all the backbones by recommending the most informative variables for observation.

5.3.3 Hyper-parameter Optimization. The proposed RAS have four important hyper-parameter: weight α in Equation 19, step size s_{adv} , step n_{adv} , learning rate lr in Algorithm 1. We use grid-search to find the best parameter (with active sensing ratio equal to 8%). Table 5 displays the searching space and the optimal values used in the training process.

6 DEPLOYMENT

Based on the sepsis prediction model and active sensing algorithm, we implement a system SepsisLab. Figure 8 and Figure 9 shows how the system is deployed in the Epic EHR Systems⁶ at OSUWMC.

SepsisLab starts to collect patients' data after the patients arrive hospital and automatically predicts sepsis risks hourly. Figure 8(A) displays a list of patients with different sepsis risk prediction scores, colored from no risk as Green, to medium risk as Yellow, to high risk as Red. When picking a patient's data, Figure 8(B) shows the patient's demographics and the dashboard that includes the patient's vital signs, lab test results, and medical history, which are helpful for clinicians to understand the patient's health states. Figure 8(C) shows the patient's sepsis risk (solid line) and uncertainty range

⁶<https://www.epic.com/software/>

Table 3: AUROC of risk prediction with the active sensing (cold-start).

	MIMIC-III				AmsterdamUMCdb				OSUWMC			
	2%	4%	6%	8%	2%	4%	6%	8%	2%	4%	6%	8%
Random Sensing	0.761	0.772	0.779	0.785	0.772	0.781	0.788	0.793	0.785	0.794	0.805	0.811
Monte Carlo	0.771	0.789	0.797	0.812	0.782	0.795	0.802	0.817	0.797	0.818	0.855	0.886
Active Learning	0.773	0.791	0.804	0.817	0.780	0.800	0.805	0.816	0.802	0.820	0.857	0.889
VAT	0.783	0.801	0.812	0.822	0.786	0.802	0.815	0.823	0.809	0.844	0.889	0.916
RAS^N	0.770	0.788	0.796	0.810	0.780	0.793	0.802	0.814	0.795	0.816	0.853	0.881
RAS^L	0.783	0.801	0.812	0.824	0.785	0.801	0.818	0.822	0.814	0.848	0.877	0.917
RAS (Ours)	0.792	0.810	0.823	0.835	0.795	0.809	0.828	0.840	0.823	0.857	0.889	0.929

Table 4: Sepsis prediction performance of three versions of RAS on observational data.

Method	MIMIC-III	AmsterdamUMCdb	OSUWMC
RAS^N	0.820	0.820	0.903
RAS^L	0.832	0.834	0.925
RAS	0.837	0.849	0.934


Table 5: The search space of hyper-parameters and optimal parameters utilized during the model training.

Parameters	Search Space	Optimal Value
Weight α	[0.1, 0.3, 0.5, 0.7, 0.9]	0.5
Learning rate lr	[1e-3, 1e-4, 1e-5]	1e-4
Step size s_{adv}	[1e-2, 1e-3, 1e-4, 1e-5]	1e-3
Step n_{adv}	[1, 2, 5, 10, 15, 20]	15

(gray area) at different times and an actionable lab item test recommendation list from SepsisLab. The items are ranked by their importance to reduce the uncertainty of the sepsis future prediction.

The interactive process with our system is visualized in Figure 9. This UI currently illustrates that a clinical expert is examining a high-risk patient’s data who was admitted 4 hours ago. The SepsisLab suggests the expert collect more lab results. The expert is interacting with the visualization to see if Lactate and Creatinine lab results were added, and how the sepsis prediction and its uncertainty would change. The clinician can select a lab item (Figure 9(b)) or multiple lab items (Figure 9(c)) and see the expected influence of the lab test result on the model uncertainty via a counterfactual prediction. By comparing different combinations of the lab test items, the clinician can obtain a better understanding of the model and make the decision to order appropriate lab tests to collect the actual item values, which then truly update the model’s prediction trajectory and uncertainty range.

Note that we used OSUWMC data for our algorithm illustration. All patients’ names and demographic info in this Figure 8 are randomly generated for illustration purposes. Ongoing deployment also includes recruit clinicians for usability evaluation to quantitative and qualitatively measure clinical outcome and user satisfaction of SepsisLab (OSUWMC IRB#: 2020H0018).

**Figure 9: The Interactive Lab Test Recommendation Module in SepsisLab System.**

7 CONCLUSION

In this work, we study a real-world problem that how to accurately predict sepsis with limited variables available. Missing values widely exist in clinical data and can cause inaccurate prediction and high uncertainty for the sepsis prediction models. To the best of our knowledge, it is the first work that studies the model uncertainty caused by missing values. We define a new term propagated uncertainty to describe the uncertainty, which is the downstream models’ uncertainty propagated from the uncertain input (i.e., imputation results). We further propose uncertainty propagation methods to quantify the propagated uncertainty. Based on the uncertainty quantification, we propose a robust active sensing algorithm to reduce the uncertainty by actively recommending clinicians to observe the most informative variables. The experimental results on real-world datasets show that the introduced propagated uncertainty is dominant at the beginning of patients’ admissions to the hospital due to the very limited variables and the proposed active sensing algorithm can significantly reduce the propagated uncertainty and thus improve the sepsis prediction performance. Finally, we design a SepsisLab system for deployment to integrate into clinicians’ workflow, which paves the way for human-AI collaboration and early intervention for sepsis management.

8 ACKNOWLEDGMENTS

This work was funded in part by the National Science Foundation under award number IIS-2145625, by the National Institutes of Health under award number R01GM141279 and R01AI188576, and by The Ohio State University President’s Research Excellence Accelerator Grant.

REFERENCES

- [1] Moloud Abdar, Farhad Pourpanah, Sadiq Hussain, Dana Rezadegan, Li Liu, Mohammad Ghavamzadeh, Paul Fieguth, Xiaochun Cao, Abbas Khosravi, U Rajendra Acharya, et al. 2021. A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Information Fusion* 76 (2021), 243–297.
- [2] Inci M Baytas, Cao Xiao, Xi Zhang, Fei Wang, Anil K Jain, and Jiayu Zhou. 2017. Patient subtyping via time-aware LSTM networks. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*. 65–74.
- [3] Roger C Bone, Robert A Balk, Frank B Cerra, R Phillip Dellinger, Alan M Fein, William A Knaus, Roland MH Schein, and William J Sibbald. 1992. Definitions for sepsis and organ failure and guidelines for the use of innovative therapies in sepsis. *Chest* 101, 6 (1992), 1644–1655.
- [4] Armen Der Kiureghian and Ove Ditlevsen. 2009. Aleatory or epistemic? Does it matter? *Structural safety* 31, 2 (2009), 105–112.
- [5] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [6] Stephen C Hora. 1996. Aleatory and epistemic uncertainty in probability elicitation with an example from hazardous waste management. *Reliability Engineering & System Safety* 54, 2-3 (1996), 217–223.
- [7] Md Mohaimenul Islam, Tahmina Nasrin, Bruno Andreas Walther, Chieh-Chen Wu, Hsuan-Chia Yang, and Yu-Chuan Li. 2019. Prediction of sepsis patients using machine learning approach: a meta-analysis. *Computer methods and programs in biomedicine* 170 (2019), 1–9.
- [8] Daniel Jarrett and Mihaela Van Der Schaar. 2020. Inverse active sensing: Modeling and understanding timely decision-making. *arXiv preprint arXiv:2006.14141* (2020).
- [9] Alistair E.W. Johnson, Tom J. Pollard, Lu Shen, et al. 2016. MIMIC-III, a freely accessible critical care database. (2016).
- [10] Sundreen Asad Kamal, Changchang Yin, Buyue Qian, and Ping Zhang. 2020. An interpretable risk prediction model for healthcare with pattern attention. *BMC Medical Informatics and Decision Making* 20 (2020), 1–10.
- [11] Alex Kendall and Yarin Gal. 2017. What uncertainties do we need in bayesian deep learning for computer vision? *Advances in neural information processing systems* 30 (2017).
- [12] James Kirchner. 2001. Data analysis toolkit# 5: uncertainty analysis and error propagation. University of California Berkeley Seismological Laboratory. Available online at: http://seismo.berkeley.edu/~kirchner/eps_120/Toolkits/Toolkit_05.pdf (2001).
- [13] Vincent Liu, Gabriel J. Escobar, John D. Greene, et al. 2014. Hospital Deaths in Patients With Sepsis From 2 Independent Cohorts. *JAMA* 312, 1 (07 2014), 90–92.
- [14] Vincent X Liu, Vikram Fielding-Singh, John D Greene, Jennifer M Baker, Theodore J Iwashyna, Jay Bhattacharya, and Gabriel J Escobar. 2017. The timing of early antibiotics and hospital mortality in sepsis. *American journal of respiratory and critical care medicine* 196, 7 (2017), 856–863.
- [15] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 2017. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083* (2017).
- [16] Lu Men, Noyan Ilk, Xinlin Tang, and Yuan Liu. 2021. Multi-disease prediction using LSTM recurrent neural networks. *Expert Systems with Applications* 177 (2021), 114905.
- [17] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. 2018. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence* 41, 8 (2018), 1979–1993.
- [18] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Jonathan Uesato, and Pascal Frossard. 2019. Robustness via curvature regularization, and vice versa. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9078–9086.
- [19] Chongli Qin, James Martens, Sven Gowal, Dilip Krishnan, Krishnamurthy Dvijotham, Alhussein Fawzi, Soham De, Robert Stanforth, and Pushmeet Kohli. 2019. Adversarial robustness through local linearization. *Advances in Neural Information Processing Systems* 32 (2019).
- [20] Kui Ren, Tianhang Zheng, Zhan Qin, and Xue Liu. 2020. Adversarial attacks and defenses in deep learning. *Engineering* 6, 3 (2020), 346–360.
- [21] Matthew A Reyna, Christopher S Josef, Russell Jeter, Supreeth P Shashikumar, M Brandon Westover, Shamim Nemati, Gari D Clifford, and Ashish Sharma. 2019. Early prediction of sepsis from clinical data: the PhysioNet/Computing in Cardiology Challenge 2019. *Critical Care Medicine* (2019).
- [22] Halden F Scott, Emily E Greenwald, Lalit Bajaj, Sara J Deakyn Davies, Lina Brou, and Allison Kempe. 2018. The sensitivity of clinician diagnosis of sepsis in tertiary and community-based emergency settings. *The Journal of Pediatrics* 195 (2018), 220–227.
- [23] Robin Senge, Stefan Bösner, Krzysztof Dembczyński, Jörg Haasenritter, Oliver Hirsch, Norbert Donner-Banzhoff, and Eyke Hüllermeier. 2014. Reliable classification: Learning classifiers that distinguish aleatoric and epistemic uncertainty. *Information Sciences* 255 (2014), 16–29.
- [24] Mervyn Singer, Clifford S Deutschman, Christopher Warren Seymour, et al. 2016. The third international consensus definitions for sepsis and septic shock (Sepsis-3). *Jama* 315, 8 (2016), 801–810.
- [25] Mervyn Singer, Clifford S Deutschman, Christopher Warren Seymour, Manu Shankar-Hari, Djillali Annane, Michael Bauer, Rinaldo Bellomo, Gordon R Bernard, Jean-Daniel Chiche, Craig M Coopersmith, et al. 2016. The third international consensus definitions for sepsis and septic shock (Sepsis-3). *Jama* 315, 8 (2016), 801–810.
- [26] Gary B Smith, David R Prytherch, Paul Meredith, Paul E Schmidt, and Peter I Featherstone. 2013. The ability of the National Early Warning Score (NEWS) to discriminate patients at risk of early cardiac arrest, unanticipated intensive care unit admission, and death. *Resuscitation* 84, 4 (2013), 465–470.
- [27] Christian P Subbe, M Kruger, Peter Rutherford, and L Gemmel. 2001. Validation of a modified Early Warning Score in medical admissions. *Qjm* 94, 10 (2001), 521–526.
- [28] Laura P Swiler, Thomas L Paez, and Randall L Mayes. 2009. Epistemic uncertainty quantification tutorial. In *Proceedings of the 27th International Modal Analysis Conference*.
- [29] Patrick Thorral, Jan Peppink, Ronald Driessen, et al. 2020. AmsterdamUMCdb: The First Freely Accessible European Intensive Care Database from the ESICM Data Sharing Initiative. (2020). <https://doi.org/10.1109/JBHI.2020.2995139> access: <https://www.amsterdammedicaldatascience.nl>.
- [30] Alexander Tsertsadze, Pam Royle, and Noel McCarthy. 2015. Community-onset sepsis and its public health burden: protocol of a systematic review. *Systematic reviews* 4, 1 (2015), 1–13.
- [31] Zhenxing Xu, Jingyuan Chou, et al. 2019. Identifying Sub-Phenotypes of Acute Kidney Injury using Structured and Unstructured Electronic Health Record Data with Memory Networks. *arXiv preprint arXiv:1904.04990* (2019).
- [32] Chao Yan, Cheng Gao, Xinneng Zhang, You Chen, and Bradley Malin. 2019. Deep imputation of temporal data. In *2019 IEEE International Conference on Healthcare Informatics (ICHI)*. IEEE, 1–3.
- [33] Changchang Yin, Ruqi Liu, Dongdong Zhang, and Ping Zhang. 2020. Identifying sepsis subphenotypes via time-aware multi-modal auto-encoder. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*. 862–872.
- [34] Changchang Yin, Rongjian Zhao, Buyue Qian, Xin Lv, and Ping Zhang. 2019. Domain Knowledge guided deep learning with electronic health records. In *2019 IEEE International Conference on Data Mining (ICDM)*. IEEE, 738–747.
- [35] Jinsung Yoon, James Jordon, and Mihaela Schaar. 2019. ASAC: Active sensing using Actor-Critic models. In *Machine Learning for Healthcare Conference*. PMLR, 451–473.
- [36] Jinsung Yoon, William R Zame, and Mihaela Van Der Schaar. 2018. Deep sensing: Active sensing using multi-directional recurrent neural networks. In *International Conference on Learning Representations*.
- [37] Shipeng Yu, Balaji Krishnapuram, Romer Rosales, and R Bharat Rao. 2009. Active sensing. In *Artificial Intelligence and Statistics*. PMLR, 639–646.
- [38] Dongdong Zhang, Changchang Yin, Katherine M Hunold, Xiaoqian Jiang, Jeffrey M Caterino, and Ping Zhang. 2021. An interpretable deep-learning model for early prediction of sepsis in the emergency department. *Patterns* 2, 2 (2021).

A APPENDIX

A.1 Uncertainty Propagation for Linear Function

For a linear function $f_w(x) = \sum_i w_i x_i$ ($1 \leq i \leq n$), the uncertainty is defined as the variance:

$$\begin{aligned}
 Var(f_w(x)) &= \int_x (f_w(x) - f_w(\mu_x))^2 dx \\
 &= \int_x \left(\sum_i w_i x_i - \sum_i w_i \mu_i \right)^2 dx \\
 &= \int_x \left[\sum_i (w_i x_i - w_i \mu_i) \right]^2 dx \\
 &= \int_x \sum_i (w_i x_i - w_i \mu_i) \sum_j (w_j x_j - w_j \mu_j) dx \\
 &= \int_x \sum_i \sum_j (w_i x_i - w_i \mu_i) (w_j x_j - w_j \mu_j) dx \\
 &= \int_x \sum_i \sum_{j=i} (w_i x_i - w_i \mu_i) (w_i x_i - w_i \mu_i)
 \end{aligned}$$

Table 6: AUROC of risk prediction of the propose RAS with different backbones.

	MIMIC-III				AmsterdamUMCdb				OSUWMC			
	2%	4%	6%	8%	2%	4%	6%	8%	2%	4%	6%	8%
FC	0.760	0.772	0.782	0.792	0.772	0.785	0.792	0.801	0.785	0.801	0.823	0.841
RNN	0.782	0.790	0.810	0.821	0.787	0.799	0.816	0.825	0.814	0.842	0.873	0.910
GRU	0.789	0.799	0.822	0.833	0.792	0.806	0.826	0.837	0.821	0.855	0.887	0.928
LSTM	0.792	0.81	0.823	0.835	0.795	0.809	0.828	0.840	0.823	0.857	0.889	0.929

$$\begin{aligned}
& + \sum_i \sum_{j \neq i} (w_i x_i - w_i \mu_i)(w_j x_j - w_j \mu_j) dx \\
& = \int_x \sum_i w_i^2 \sigma_i^2 + \sum_i \sum_{j \neq i} w_i w_j \rho_{i,j} \sigma_i \sigma_j dx
\end{aligned}$$

where μ_x denotes the mean of variable x . i and j denote the indices of variables or parameters. ρ denotes the correlation coefficient. σ denotes the standard deviation.

A.2 Missing Rates of Clinical Variables

We display the missing rates of lab test variables in Table 7.

Table 7: Missing rates of observed lab tests.

variable	AmsterdamUMCdb	OSUWMC	MIMIC-III
WBC	67%	78%	69%
BUN	63%	76%	66%
GCS	29%	50%	33%
Urineoutput	23%	39%	33%
Creatinine (CRT)	75%	85%	80%
Platelet (PLT)	76%	88%	82%
Glucose (GLC)	34%	49%	36%
Sodium (SDM)	55%	72%	65%
Hemoglobin (HMG)	56%	75%	69%
Chloride (CLR)	62%	70%	66%
Bicarbonate (BCB)	69%	74%	67%
Lactate (LCT)	88%	90%	89%
INR	78%	84%	80%
PTT	76%	83%	79%
Magnesium	66%	76%	69%
Aniongap (AG)	62%	78%	67%
Hematocrit (HMT)	60%	76%	64%
PT	78%	92%	80%

A.3 Model Performance with different backbones

Our model is applicable to various models, including LSTM, GRU, and fully-connected networks (FC). LSTM has shown superior performance in modeling clinical time series data in multiple tasks, including missing value imputation [32, 34], clinical prediction [16], and patient subtyping [2], so we choose LSTM as the model backbone. We also conducted more experiments with different backbones as shown in Table 6. The experimental results show that the proposed model can significantly improve the prediction performance for all the backbones by recommending the most informative variables for observation.