

NPRL: Nightly Profile Representation Learning for Early Sepsis Onset Prediction in ICU Trauma Patients

Tucker Stewart¹ Katherine Stern² Grant O’Keefe² Ankur Teredesai^{1,3} Juhua Hu^{1†}

¹ School of Engineering and Technology, University of Washington, Tacoma, USA

² Department of Surgery, University of Washington, Seattle, USA

³ CueZen Inc., Seattle, USA

{trstew, kstern, gokeefe, ankurt, juhuah}@uw.edu

Abstract—Sepsis is a syndrome that develops in the body in response to the presence of an infection. Characterized by severe organ dysfunction, sepsis is one of the leading causes of mortality in Intensive Care Units (ICUs) worldwide. These complications can be reduced through early application of antibiotics. Hence, the ability to anticipate the onset of sepsis early is crucial to the survival and well-being of patients. Current machine learning algorithms deployed inside medical infrastructures have demonstrated poor performance and are insufficient for anticipating sepsis onset early. Recently, deep learning methodologies have been proposed to predict sepsis, but some fail to capture the time of onset (e.g., classifying patients’ entire visits as developing sepsis or not) and others are unrealistic for deployment in clinical settings (e.g., creating training instances using a fixed time to onset, where the time of onset needs to be known apriori). In this paper, we first propose a novel but realistic prediction framework that predicts each morning whether sepsis onset will occur within the next 24 hours using the most recent data collected the previous night, when patient-provider ratios are higher due to cross-coverage resulting in limited observation to each patient. However, as we increase the prediction rate into daily, the number of negative instances will increase, while that of positive instances remain the same. This causes a severe class imbalance problem making it hard to capture these rare sepsis cases. To address this, we propose a nightly profile representation learning (NPRL) approach. We prove that NPRL can theoretically alleviate the rare event problem and our empirical study using data from a level-1 trauma center demonstrates the effectiveness of our proposal.

Index Terms—early sepsis onset prediction, ICU trauma patients, nightly patient profile, rare event, profile representation learning

I. INTRODUCTION

Sepsis is a syndrome of a host’s dysregulated immune response to the presence of an infection. The host immune response is not specific in its mobilization against pathogens

in the body, that it starts to attack the host’s own tissue as well, causing organ dysfunction and tissue damage. To this day, sepsis remains to be a prominent complication in the modern medical facilities, particularly ICUs. According to a global audit conducted in [1], depending on the region studied, 13.6% to 39.3% of patients admitted in ICU are affected by sepsis. Globally, this proportion is 29.5%. Of the patients that develop sepsis in the ICU, up to 26% will die likely due to sepsis. From this, it can be estimated that approximately 7.6% of patients admitted into ICUs across the globe will die as a result of sepsis, making it one of the leading causes of death in ICUs. Moreover, sepsis also contributes to long term morbidity and mortality. About one-third of ICU sepsis survivors develop persistent and prolonged organ dysfunction, a syndrome commonly recognized as chronic critical illness, characterized by persistent immune suppression [2]–[4], muscle wasting [5], and recurrent infections [6]. Survivors experience high rates of sepsis recidivism, hospital readmission, deficits in physical and cognitive function [7], and increased 1-year and 5-year all-cause mortality [8].

These outcomes can be improved through early intervention [9]. Early administration of antibiotics treats the underlying infection and prevents the progression of sepsis-related organ dysfunction. It is estimated that as much as 80% of septic shock-related deaths are preventable with early intervention and chances of survival decrease about 7.6% each hour that action is not taken [10]. Hence, it is critical to the survival and well-being of patients to anticipate the onset of sepsis accurately and early. Although patients who acquire sepsis during hospitalization (i.e., hospital-acquired sepsis) are at greatest risk for sepsis-associated morbidity and mortality [11], early detection is challenging in these groups. In critically ill trauma patients (i.e., individuals admitted to the ICU for management of injury caused by blunt or penetrating force), injury-related inflammation and organ dysfunction may increase the risk for sepsis, while also masking the clinical signs of infection [12], [13]. Therefore, detecting sepsis early in the critically ill trauma population is in great need but challenging, which is the focus of this work.

Due to the abundance of available data, many have

*Stewart and Hu’s research is supported in part by NSF (IIS-2104270). Teredesai’s research is supported in part by CueZen Inc. We would like to thank Microsoft Inc., and University of Washington eScience Institute - Azure Cloud Computing Credits for Research and Teaching grant for their support by providing us with the Cloud Compute Resources on Microsoft Azure. All opinions, findings, conclusions and recommendations in this paper are those of the author and do not necessarily reflect the views of the funding agencies.

[†]Corresponding author

looked into machine learning using patients' Electronic Health Records (EHRs) to predict sepsis or septic shock. However, many prior studies only work retrospectively [14]. Specifically, many studies like [9], [14]–[19] identify the timestamp of the target event such as sepsis onset or septic shock, and then look back in time for a fixed time interval for early prediction. As such, the time length between prediction and the target event stays fixed. This approach is less useful prospectively for deployment because clinicians often do not know in advance when sepsis will actually onset. Therefore, in practice, it would be difficult to decide how and when to use such ML models for prediction. For example, we may need to use the model very frequently (e.g., every hour), so that we do not miss any chance of early prediction. This is not realistic in a live setting. In addition, some prior approaches like [9], [14], [16], [18] created one data sample per hospital admission, not multiple instances such as for each day the patient is in the hospital. In these cases, the model does not discriminate between days, where sepsis is or is not present in the patient. Thus, they do not capture the time of onset as well, which is a major shortcoming we address in our work.

In this paper, we propose a novel but more realistic prediction setup for deployment in hospitals. Specifically, we use data from each night when patients have limited observations from doctors and re-assess the potential for sepsis onset the following day. In this manner, we position the classifier as a diagnostic tool similarly as lab tests rather than an auxiliary alarm system. Then, instead of labeling the entire visit to have sepsis (i.e., between 13.6% and 39.3% in the ICU) or not, each day within the visit is examined for the first potential occurrence of sepsis onset. Time prior to sepsis onset is treated as time at risk. However, sepsis examples become rare compared to negative examples, resulting in a severe class imbalance problem (e.g., $< 2\%$ sepsis cases in the level-1 trauma center data). Without any help to address this rare event problem, machine learning models can be easily misled to predict every example to be negative, which is useless.

In previous sepsis prediction studies [16], [18], [20]–[22], the class imbalance problem was often addressed by simply re-sampling the training data. For example, random oversampling achieves balance by repeating minority examples, but often leads to overfitting those examples in the minor class. On the other hand, random undersampling removes major examples, but results in information loss. Later, [19] proposed to use ensemble to learn various characteristics from data to alleviate the class imbalance problem, which however is very expensive. More importantly, our class imbalance problem is way more serious. Recently, due to the development of deep neural networks, deep learning has outperformed in various aspects. In terms of sepsis prediction, deep learning architectures such as Recurrent Neural Networks (RNNs) [9], [17], [19], [21] and Temporal Convolutional Neural Networks (TCNs) [18], [20], which can capture temporal and sequential patterns significantly outperformed traditional machine learning methods. Moreover, deep learning has also helped advance techniques to address the class imbalance problem [23]. For example, in

computer vision, [24] proposed a balanced loss that weights the loss for a particular class inversely proportional to number of instances in that class and [25] showed that self-supervised representations are more robust to class imbalance than supervised representations. However, this has been rarely studied in early sepsis onset prediction, where EHR data are totally different from images in computer vision.

Therefore, to do nightly sepsis prediction resulting in rare sepsis events, we propose to do self-supervised Nightly patient Profile Representation Learning (NPRL). During the representation learning, we aim to capture the unique perspectives of different patients at different nights. We theoretically prove that such representations learned from self-supervised learning can preserve the diversity of patient profiles at different nights, and thus examples from the minor class will not be dominated by those from the major class. Our empirical study conducted on traumatic patients from a level-1 trauma center demonstrates the effectiveness of our proposal. The main contributions of this work can be summarized as follows.

- We propose a novel nightly prediction setup for sepsis in ICU trauma patients that is more deploy-able in prospective clinical environments of ICUs and meets the needs of ICU staff.
- We develop a self-supervised nightly patient profile representation learning method to capture unique perspectives of different patients at different nights, as to alleviate the severe class imbalance problem with a theoretical guarantee.
- We demonstrate the effectiveness of our proposal using data from a level-1 trauma center and compare it with various existing state-of-the-art methods.

II. RELATED WORK

In this section, we briefly review different sepsis prediction setups, machine learning methodologies used in the literature for early sepsis prediction, and existing techniques applied to address its class imbalance problem.

A. Historical Sepsis Prediction Setup

How the prediction setup for sepsis is framed has many major consequences for the utility of the predicted label and the model fitting, which has been discussed in depth by Lauritsen et al. [26]. In the reviewed literature, there are three different setups used, that is, fixed time to onset, sequential, and sliding window.

Most studies used a fixed time to onset framing for their prediction setup [14]–[19]. For example, [16] predicted onset in 12, 24, or 48 hours prior to onset. However, as mentioned by [9], this setup is learning to predict onset at fixed intervals prior to onset, and thus they are often limited to these retrospective studies, where the time of onset is known a priori. Thereafter, in deployment, these models will need to assess patients' risk at regular interval as not to miss the chance to capture onset. The next most used prediction setup is sequential [20]–[22], [27], which generates predictions at regular intervals using all the data available from admission to

the time of prediction. This is more realistic in that it expects to assess patients at regular intervals and it considers all relevant information available. However, a major challenge is the variable observation window. Moreover, the observation window grows the longer the patient stays in the hospital. Thereafter, the amount of data to include can be quite large, and thus the data is usually aggregated to distill the information down.

Sliding window [28]–[30] is similar to sequential, except that it uses a fixed time length prior to prediction for the observation window, making the shape of the input fixed and only considering the most recent data. With this approach, the raw EHR data can be used directly, preserving the temporally progressing patterns. Hence, sliding window approach is preferable for our purpose in discerning when the onset of sepsis occurs. However, there are unmet needs in medical infrastructures. For example, during night hours, the patient-provider ratios are much higher due to cross-coverage, and thus there is limited observations to each patient. An assistant prediction from machine learning in situations when we do not have full observations of each patient would be valuable. In this work, we focus on whether sepsis onset will occur within a 24-hour window with the help of most recent data collected at night. We time delivery of daily predictions for the early morning, right before routine morning rounds, so that medical staff can apply this information to patient care decisions throughout the day.

B. Machine Learning for Sepsis Prediction

In this subsection, we delve into the machine learning technologies used to predict sepsis and that used to address its class imbalance problem.

1) *Sepsis Prediction*: In the literature, various traditional machine learning methods have been adopted for sepsis prediction including but not limited to support vector machines [31], random forest [16], and XGBoost [32]. A well-known sepsis early warning system currently deployed in hospitals is Epic Sepsis Model (ESM). ESM is a penalized logistic regression model that produces a score indicating the potential of sepsis calculated every fifteen minutes. Recently, [33] found that ESM performed poorly using data from Michigan Medicine. While there is still some utility of ESM scores to help medical decisions, there is a need for better early predictions of sepsis onset to be deployed in medical infrastructures. It should be noted that we aim to move away from early warning systems that produce frequent false alarms frustrating medical staff.

Recently, due to the advances in deep neural networks and their utility of automatic feature learning, deep learning methods have been increasingly applied to sepsis related prediction. For instance, [15], [27] transformed the raw time-series data through a Multitask Gaussian Process (MGP), and then fed the latent function values through an RNN for the classification, while [20] extended that to apply the Temporal Convolutional neural Network (TCN). [21] also employed a sequential approach using patients' vital signs, demographics, lab results, and medications to predict sepsis onset for the next day. The selected features are aggregated daily so that

each day, up to 14 days, is one time step in the RNN. [18], [22] used an event-based data representation, where each event is formatted as a vector with three values (i.e., timestamp, name of event, and the value) and CNN-RNN or TCN are adopted for sepsis prediction. In this work, we focus more on addressing the severe class imbalance problem. Considering the prevalent promising performance of RNN in capturing temporal patterns for sepsis prediction, we adopt it as our partial deep learning architecture for nightly patient profile representation learning.

2) *Class Imbalance for Sepsis*: Most existing sepsis prediction studies [16], [18], [20]–[22], [34], [35] address the class imbalance problem in sepsis prediction by simply resampling the data to reduce the class imbalance. However, simple over or under sampling strategies do not solve the problem well [36]. Random oversampling achieves balance by repeating minority examples, often leading to overfitting on those examples in the minority class, and random under-sampling removes majority examples sacrificing much of the available information. To better address the class imbalance problem, [19] utilizes an ensemble of Long Short-Term Memory networks (LSTM). Specifically, multiple LSTM models are trained, capturing different characteristics of the data. Then, their predictions do ensemble to better capture sepsis prediction, which is obviously expensive.

It should be noted that the development of deep learning has also advanced the machine learning techniques to address the class imbalance problem [23] in other domains but not sepsis. For example, in computer vision, [24] proposed a balanced loss that weights the loss for a particular class inversely proportional to number of instances in that class and [25] showed that self-supervised representations are more robust to class imbalance than supervised representations. However, this has been rarely studied in early sepsis onset prediction, where EHR data are totally different from images in computer vision. Based on our proposed prediction setup, we develop a nightly patient profile representation learning using EHR data to address the serious class imbalance problem.

III. NIGHTLY PROFILE REPRESENTATION LEARNING

In this section, we first describe our nightly prediction setup for early sepsis prediction, which is more realistic in ICUs. Then, due to the resulting severe class imbalance problem making sepsis a rare event, we propose a patient profile representation learning method to capture the unique perspective of different patients at different nights.

A. Early Sepsis Prediction Setup

In this study, we propose to utilize data collected from patients at night to predict whether the onset of sepsis will occur within the next 24 hours. Specifically, data recorded between the hours of 10 p.m. and 6 a.m. are used to predict whether sepsis onset will occur in the next 24 hours (i.e., until 6 a.m. of the next day). We use data recorded during this nighttime window because 1) most staff are gone for the night so the model captures a gap in observation, 2) data collected




Fig. 1. Sepsis Prediction Setup. Temporal features are extracted from the nighttime window between 10 p.m. and 6 a.m., forming a $9 \times T$ matrix. S static features which are consistent throughout the admission for each patient. Both are used for predicting sepsis onset within the subsequent 24 hour prediction window.

at night is more likely to reflect the actual physiology of the patient as they are less exposed to external stimuli of hospital staff, 3) there are fewer interruptions to data collection as diagnostic procedures and interventions are typically planned during daytime hours, and 4) the night window is immediately adjacent to the period of time in which we predict sepsis onset. Each day in a typical ICU, teams of medical staff conduct morning rounds, a standard procedure in which staff will evaluate the current progress of patients and plan treatment for the next 24 hours. This is why we target predictions for early morning. This way, the deployed classifier will act as another decision support diagnostic tool similar to lab tests for staff to discern during rounds. This is much more desirable to clinicians compared to auxiliary monitoring systems that lead to alarm fatigue. Therefore, our setup is more practical and complements current ICU operating procedures.

It should be noted that for each patient in ICUs, we are able to collect hourly vital signs (e.g., heart rate and body temperature) and hourly cumulative exposures (e.g., IV fluid bolus volume). Given this kind of hourly temporal data, our sepsis prediction problem can be framed as a multi-variate time-series classification task, which is beneficial to adopt RNN to capture temporal patterns for sepsis prediction. Some patient profile information (e.g., age, gender, and injury type) could be potential risk factors contributing to sepsis, but are static. Therefore, in our setting each observation contains 9 hours (i.e., 10 p.m. to 6 a.m.) of T temporal feature values and S static features as illustrated in Fig. 1. Then, we aim to predict if sepsis will be onset in the next 24 hours from 6 a.m. of the current day to 6 a.m. the next morning as shown in Fig. 1.

B. NPRL for Class Imbalance

Given both the temporal and static data for each observation in our setting, it is natural to adopt a multi-modal RNN for sepsis prediction as illustrated in Fig. 2, which is denoted as

model f_θ with parameters θ . Concretely, the $9 \times T$ matrix of temporal data is fed to a bidirectional RNN layer to capture the temporal information. LSTM [37] and Gated Recurrent Unit (GRU) [38] are often used, where GRU has been proved to be preferable in medical prediction tasks [39]. Moreover, bidirectional RNNs can further improve the performance [40]. Therefore, we opt for a Bidirectional GRU layer (i.e., the blue component in Fig. 2) for our temporal representation learning. Then, the representation learned from the temporal features is concatenated with the representation learned from the S static features using a deep dense layer architecture (i.e., the red component in Fig. 2) for the final classification of sepsis.

Let $\{(x_i^T, x_i^S), y_i\}_{i=1}^n$ denote our training data set, where x_i^T denotes the temporal data, x_i^S denotes the static data, and $y_i \in \{1, \dots, C\}$ is the corresponding label with $C = 2$ for sepsis. After feeding our temporal and static data to the multi-modal RNN model f_θ , we can denote the representation in the last fully connected layer before classification as $\mathbf{x}_i = f_\theta(x_i^T, x_i^S)$ (i.e., the purple representation in Fig. 2). Then, the classification model can be learned by empirical risk minimization (ERM) as

$$\min_{\theta} \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{x}_i, y_i; \theta)$$

where $\ell(\cdot)$ is a loss function and cross-entropy loss is popular in deep learning. For a C -class classification problem, the task is equivalent to

$$\min_{\theta} \sum_{c=1}^C \frac{n_c}{n} g(\mathbf{x}, y, c) \quad (1)$$

where n_c shows the number of examples from the c -th class and

$$g(\mathbf{x}, y, c) = \frac{1}{n_c} \sum_{i: y_i=c} \ell(\mathbf{x}_i, y_i; \theta)$$

According to the formulation, it is obvious that the loss will be dominated by the major class if there is some class j such that




Fig. 2. Multi-Modal RNN. Temporal data is processed using a bidirectional GRU layer (blue) to capture temporal information while static data is processed using a deep dense layer architecture (red). Results from each are combined via concatenation for further dense layers (purple).

$\forall c, c \neq j, n_j \gg n_c$. For example, if we have 100 examples for sepsis prediction with only 2 of them having sepsis, the loss will be dominated by those 98 non-sepsis examples, since we treat each example equivalently. Then, the model may directly predict each example as non-sepsis, which is useless.

Fortunately, in the following, we can show that the self-supervised representation learning can help alleviate the problem. Before we do the sepsis classification directly, we can conduct a self-supervised learning on f_θ without the sepsis classification layer to initialize the model's parameters as θ_0 . Specifically, given $\{(x_i^T, x_i^S)\}_{i=1}^n$ without sepsis assignments, we aim to uniquely identify each observation in the training data, which is corresponding to a night profile for a patient at a specific night and thus named as Nightly Profile Representation Learning (NPRL). Thereafter, θ_0 can be learned by setting the classification layer as a n -class classification problem with training data as $\{(x_i^T, x_i^S), i\}_{i=1}^n$, where each instance is uniquely identified by the label i .

Given f_{θ_0} , after changing the classification to sepsis, the sepsis classification model can be fine-tuned with a small learning rate and the optimization problem becomes

$$\min_{\theta} \sum_{c=1}^C \frac{n_c}{n} g(\mathbf{x}, y, c) \quad \text{s.t.} \quad \|\theta - \theta_0\|_F \leq \gamma \quad (2)$$

which is equivalent to

$$\min_{\theta} \sum_{c=1}^C \frac{n_c}{n} g(\mathbf{x}, y, c) + \frac{\lambda}{2} \|\theta - \theta_0\|_F^2 \quad (3)$$

Compared with the original objective in Eqn. 1, a regularization term is introduced to constrain the difference between the learned model and the initialized model by NPRL. The benefit from NPRL initialized model can be demonstrated in the following theorem.

Definition 1. [41] A function h is L -Lipschitz continuous if

$$\|h(x) - h(y)\|_2 \leq L\|x - y\|_2$$

For the sake of convenience, we denote $\mathbf{x} = \theta(x)$ as a simplified definition of $\mathbf{x} = f_\theta(x^T, x^S)$.

Theorem 1. Let $\theta(x)$ be L -Lipschitz in x following [42] and θ^* denote the solution of Eqn. 3. If $\gamma \leq \frac{1}{8L}$, we have

$$\begin{aligned} \forall i, j, \quad & \|\theta^*(x_i) - \theta^*(x_j)\|_2^2 \\ & \geq \|\theta_0(x_i) - \theta_0(x_j)\|_2^2 - \|\theta_0(x_i) - \theta_0(x_j)\|_2/2 - 1/32 \end{aligned}$$

Proof.

$$\begin{aligned} & \|\theta^*(x_i) - \theta^*(x_j)\|_2^2 \\ &= \|\theta^*(x_i) - \theta_0(x_i) + \theta_0(x_j) - \theta^*(x_j) + \theta_0(x_i) - \theta_0(x_j)\|_2^2 \\ &= \|\theta_0(x_i) - \theta_0(x_j)\|_2^2 + 2\langle \theta^*(x_i) - \theta_0(x_i), \theta_0(x_j) - \theta^*(x_j) \rangle \\ &+ \|\theta^*(x_i) - \theta_0(x_i)\|_2^2 + 2\langle \theta^*(x_i) - \theta_0(x_i), \theta_0(x_i) - \theta_0(x_j) \rangle \\ &+ \|\theta_0(x_j) - \theta^*(x_j)\|_2^2 + 2\langle \theta_0(x_j) - \theta^*(x_j), \theta_0(x_i) - \theta_0(x_j) \rangle \\ &\geq \|\theta_0(x_i) - \theta_0(x_j)\|_2^2 - 2\|\theta^*(x_i) - \theta_0(x_i)\|_2 \|\theta_0(x_j) - \theta^*(x_j)\|_2 \\ &- 2\|\theta^*(x_i) - \theta_0(x_i)\|_2 \|\theta_0(x_i) - \theta_0(x_j)\|_2 \\ &- 2\|\theta_0(x_j) - \theta^*(x_j)\|_2 \|\theta_0(x_i) - \theta_0(x_j)\|_2 \end{aligned}$$

Due to the smoothness of $\theta(x)$, we have

$$\begin{aligned} & \|\theta^*(x_i) - \theta^*(x_j)\|_2^2 \\ & \geq \|\theta_0(x_i) - \theta_0(x_j)\|_2^2 - 2L^2\gamma^2 - 4L\gamma\|\theta_0(x_i) - \theta_0(x_j)\|_2 \end{aligned}$$

We can easily observe the result with $\gamma \leq \frac{1}{8L}$. $\square$

Remark: The above theorem illustrates that fine-tuning based on a pre-trained model appropriately can preserve the diversity in the pre-trained representations, which can avoid the collapse of minor classes in a class imbalance problem. In our pre-trained model, we aim to uniquely identify each night of each patient. Therefore, for all night instances, we have

$$E_{\forall i, j \in \{1, \dots, n\}, i \neq j} [\theta_0(x_i)^\top \theta_0(x_j)] = 0$$

With this observation from NPRL to uniquely identify each night of each patient, we have the bound as follows.

Corollary 1. *Let $\theta(x)$ be L -Lipschitz in x and θ^* denote the solution of Eqn. 3. If $\gamma \leq \frac{1}{8L}$ and θ_0 is pre-trained with NPRL and $\|\theta_0(x)\|_2 = \|\theta^*(x)\|_2 = 1$, we have*

$$E_{\forall i, j \in \{1, \dots, n\}, i \neq j} [\theta^*(x_i)^\top \theta^*(x_j)] \leq 0.37$$

Proof. Due to Jensen's inequality, we have

$$E[\|\theta_0(x_i) - \theta_0(x_j)\|_2] \leq \sqrt{E[\|\theta_0(x_i) - \theta_0(x_j)\|_2^2]} = \sqrt{2}$$

Then, we can take it back to the inequality by expanding $\|\cdot\|_2^2$ as

$$\begin{aligned} \forall i, j, \quad & \|\theta^*(x_i)\|_2^2 + \|\theta^*(x_j)\|_2^2 - 2\theta^*(x_i)^\top \theta^*(x_j) \\ & \geq \|\theta_0(x_i)\|_2^2 + \|\theta_0(x_j)\|_2^2 - 2\theta_0(x_i)^\top \theta_0(x_j) \\ & - \|\theta_0(x_i) - \theta_0(x_j)\|_2^2 / 2 - 1/32 \end{aligned}$$

By applying the expectation on both sides, we have

$$\begin{aligned} & E_{\forall i, j \in \{1, \dots, n\}, i \neq j} [\theta^*(x_i)^\top \theta^*(x_j)] \\ & \leq E_{\forall i, j \in \{1, \dots, n\}, i \neq j} [\theta_0(x_i)^\top \theta_0(x_j)] + E[\|\theta_0(x_i) - \theta_0(x_j)\|_2] / 4 \\ & + 1/64 \leq 0 + \sqrt{2}/4 + 1/64. \end{aligned}$$

Then, we can easily observe the bound. $\square$

Explicitly, after fine-tuning based on the pre-trained model by NPRL, although the representations are no longer orthogonal in expectation, the diversity can be well preserved. That means the examples from the minor class will not be dominated by those from the major class.

In summary, we propose to do nightly patient profile representation learning, that is NPRL, in which we first conduct a self-supervised representation learning to uniquely identify each night of each patient. Then, the sepsis classification model is initialized by the pre-trained model from NPRL except the classification layer. Finally, the sepsis classification model can be learned by fine-tuning the pre-trained model with a small learning rate to alleviate the class imbalance problem.

IV. EXPERIMENTS

To demonstrate the proposed method, we evaluate it on level-1 trauma center data described as follows.

A. Data Description

We were provided de-identified Electronic Health Record (EHR) data pertaining to patients aged 16 years and older admitted into the ICU of UW Medicine Harborview Medical Center following injury between the years of 2012 and 2019¹ [43]. In total, the data contains EHR data for 2,802 patients, 486 of which developed sepsis during their stay in the ICU, making up approximately 17%. Sepsis is defined according to the 2016 international guidelines, Sepsis-3 [44], as a clinically suspected infection associated with worsening of organ dysfunction. **We restricted to include only infections that were identified between hospital days 3 and 14, as this is a period of time when patients are at risk for hospital acquired infections [45] and still in the acute phase of critical illness.** Patients were labeled as having sepsis if they met the following two criteria within three days of each other.

- Suspected of having an infection. Infection confirmed with a positive culture sample or by chart review.
- Exhibits worsening organ dysfunction denoted by a 2 point or greater increase in Sequential Organ Failure Assessment (SOFA) [46] ($\Delta\text{SOFA} \geq 2$).

The time of onset is determined by the time that the positive culture sample was drawn from the patient due to it being the most specific landmark for Harborview Medical Center clinicians' concern that an infection is present or developing.

The provided EHR data includes details about patients' demographics, information about their injuries, physiological signatures recorded hourly, therapeutics administered, and comorbidities. Although information about patient comorbidities were available in the retrospective data and may contribute to sepsis risk, these data are often not ascertained until after the hospital course. Our intention was to develop a model that could easily translate into the clinical setting; so we intentionally omitted data that would not be reliably available at the point of care. From this data, we constructed the features into subsets. Each of these subsets encapsulated one category of data identified as being a contributing factor to the development of sepsis. Subset 1 is vital signs of the patient aggregated each hour, subset 2 contains patient factors and initial physiology from the first 48 hours, and subset 3 is cumulative exposures treated as events accumulated over time. Subsets 1 and 3 are temporal, while subset 2 is static. The purpose of parsing the data into subsets is that subsets can be dynamically isolated and combined to test their contribution to prediction as well as their interaction with one another. Table I summarizes the detailed feature information for each subset.

EHR being a large, heterogeneous data format is susceptible to large amounts of missingness and sparsity. Another benefit of our prediction setup is a short observation window. In this case, it is often reliable to apply Last Observation Carry Forward (LOCF) to impute missing values, except for Mean Arterial Pressure that was calculated from Systolic and Diastolic Blood Pressure as $\text{MAP} = (2\text{DBP} + \text{SBP})/3$ following

¹E-mail corresponding author for access to this dataset.

TABLE I

FEATURE SUBSETS. * INDICATES THAT THEY ARE COLLECTED FROM AND AGGREGATED OVER THE FIRST 48 HOURS OF THE PATIENT'S ADMISSION.

Subsets	Features
Subset 1	Heart Rate Diastolic Blood Pressure (DBP) Mean Arterial Pressure (MAP) Respiratory Rate Temperature Fraction of Inspired Oxygen (FiO2)
Subset 2	Age Sex Mechanism of Injury Was Transferred from Another Hospital? Has Head Injury? First Systolic Blood Pressure in ED Reverse Shock Index Max Base Deficit* Max Lactate* Total Red Blood Cell Units Transfused* Total Intravenous Crystalloid Volume (L) (Not in OR)* APACHE II Antibiotic Exposure* Number of Surgeries* Emergency Department Disposition
Subset 3	IV Fluid Bolus Volume Red Blood Cell Units Transfused Ventilator Days Number of Surgeries Surgery Duration

[47]. Thereafter, SBP is no longer used due to the redundancy, and DBP and MAP being more clinically relevant [48].

B. Instance Extraction and Inclusion

After cleansing the data and parsing the features into subsets, we extract instances from the nighttime window to train and evaluate our proposed method. For each visit, we extract the temporal features between the hours of 10 p.m. and 6 a.m. for each night the patient is in the ICU for days 3 through 14. Day 1 and 2 are excluded, since hospital acquired infections develop after 48-72 hours of hospitalization [45]. Infections before day 3 are considered to be acquired from the community or associated with healthcare; therefore it is not a problem of the trauma population. Windows after 14 days are also excluded as after 14 days, patients are no longer considered acutely ill but instead are considered to have chronic critical illness, which has different phenotype from sepsis in the acute recovery phase. In addition, instances containing any null values after LOCF are also excluded. For each visit, the static features are identified and added to each instance.

Lastly, above instances are labeled according to if the first occurrence of sepsis onset occurs within 24 hours after the observation window. While sepsis recidivism is a big problem as well, our setup targets the first sepsis event for the purpose of early prediction. Therefore, instances after the first onset for septic patients were excluded. This results in a total of 25,952 instances with 471 positive ones and 25,481 negative ones, where the serious class imbalance problem can be observed ($< 2\%$ for sepsis).

C. Experimental Setup

To fairly evaluate the performance without bias, 5-fold stratified cross validation was used such that there are five folds of approximately 5,190 instances. Multiple iterations of training and testing are done so that each fold would be used as a test fold once while the other four folds are used for training. The folds are stratified to preserve the original class distribution. In addition, due to the class imbalance between negative and positive instances, the majority class would dominate the empirical risk minimization during the training. In this case, the model would learn to always predict negative. Therefore, during the training of the sepsis vs. non-sepsis classification except for experiments with class-balanced loss, we under-sampled the negative instances and over-sampled the positive instances for the training data only, each with 2,600.

We compare the proposed method NPRL with several existing state-of-the-art methods as follows.

- XGBoost [49]: the state-of-art traditional machine learning method with all available features flattened and concatenated as a long feature vector.
- LSTM w/ Attention [21]: an existing RNN model that can do daily prediction but can only use temporal data.
- Multi-Modal RNN without NPRL: In GRU, each unit expands the representation of the T features to a 256-dimensional vector such that the output dimension of this layer is 9×512 ; The hidden state of each recurrent unit is passed further down the network to incorporate information from each hour of the observation. As such, a flattening layer is used to flatten the hidden state from the recurrent layer to form a 4608-dimensional vector so it can be combined with hidden state from the static component. Static features are processed through a deep neural network comprised of an input layer, followed by three dense layers consisting of 16, 8, and 1 hidden unit(s), respectively. The results of each component are then combined via a concatenation layer that concatenates the hidden states together to form a 4609-dimensional vector that can then be used for classification.
- Multi-Modal RNN with Class Balance Loss from [24].

For deep learning models, numerical features were each scaled to $[0, 1]$ using Min-Max scaling. Moreover, the static component of the Multi-Modal RNN is excluded from the architecture if no static features are included in the corresponding configuration thus making it just an RNN. Lastly, each of these experiments were evaluated using area under the receiver operating characteristic curve (AUROC). AUROC is often preferred for evaluating binary classification results and is the most used metric to get a holistic understanding of the performance of the classifier. However, it does not provide any information about how the model is performing for positive and negative classes separately. Therefore, we also assess confusion matrix counts along with the sensitivity and specificity. Each of these metrics are calculated for the data in the test fold for each iteration of cross validation. In order to consider all these evaluation metrics extensively, instead of

TABLE II

PERFORMANCE OF XGBOOST AND MULTI-MODAL RNN USING DIFFERENT FEATURE SUBSETS. THE RESULTS ARE BASED ON THE AGGREGATION (I.E., NOT AVERAGE) OF 5 TEST FOLDS. THE BEST IS IN BOLD AND THE 2ND BEST IS UNDERLINED.

Feature Set	AUROC	True Positives	True Negatives	Sensitivity	Specificity
XGBoost					
Subset 1	0.7061	287	17577	0.6093	0.6898
Subsets 1+2	0.7130	296	17161	0.6285	0.6735
Subsets 1+3	0.7398	315	17632	0.6688	0.6920
Subsets 1+2+3	0.7429	317	17749	0.6730	0.6966
Multi-Modal RNN					
Subset 1	0.7097	295	17710	0.6263	0.6950
Subsets 1+2	0.7100	300	17187	0.6369	0.6745
Subsets 1+3	0.7794	346	17793	0.7346	0.6983
Subsets 1+2+3	0.7716	<u>332</u>	18108	0.7049	0.7106

TABLE III

PERFORMANCE COMPARISON USING SUBSETS 1 AND 3. THE RESULTS ARE BASED ON THE AGGREGATION (I.E., NOT AVERAGE) OF 5 TEST FOLDS. THE BEST IS IN BOLD AND THE 2ND BEST IS UNDERLINED.

Model	AUROC	True Positives	True Negatives	Sensitivity	Specificity
XGBoost	0.7398	315	17632	0.6688	0.6920
LSTM w/ Attention [21]	0.7610	<u>381</u>	15222	<u>0.8089</u>	0.5974
RNN	0.7794	346	<u>17793</u>	0.7346	<u>0.6983</u>
RNN with Class Balanced Loss	0.7267	316	17480	0.6709	0.6860
RNN with Class Balanced Loss (Undersampling)	0.7733	325	18028	0.6900	0.7075
RNN with NPRL	0.7870	390	15602	0.8280	0.6120

averaging over 5 folds, we aggregated them into one result. For example, when each fold is used as the test data, we have the corresponding number of true positives. Then, the total number of true positives is counted by the summation over all five test folds and the ROC curve is the intercept of each curve. It is mainly due to the reason that averaged values over counts like true positives are hard to interpret. Our code is available here².

D. Performance Comparison

Before evaluating the efficacy of our proposed methodology, we first assessed the contribution factor each feature subset available in our data to identify an optimal combination by applying two baseline models (i.e., XGBoost and Multi-Modal RNN) to four different feature subset combinations (i.e., Subset 1 only, Subset 1 + Subset 2, Subset 1 + Subset 3, and Subset 1 + Subset 2 + Subset 3).

Table II compares the performance of using different feature subsets. First, it can be observed from the comparison between feature sets that accumulated features in subset 3 can significantly improve the performance for both models. However, the contribution of static features is limited, which may indicate that recent temporal progressing pattern during the night is more useful to predict sepsis onset. Second, RNN with the ability to capture the temporal progressing pattern outperforms XGBoost that treats each feature independently, especially for positive cases. This confirms the effectiveness of temporal models on sepsis prediction as in the literature. Without the loss of common practice, we also report the average performance over the 5 test folds with the standard

deviation in Table IV, in which the similar phenomenon as the aggregated results can be observed.

TABLE IV

AVERAGE PERFORMANCE OF XGBOOST AND MULTI-MODAL RNN USING DIFFERENT FEATURE SUBSETS. THE AVERAGE PERFORMANCE OVER 5 TEST FOLDS WITH THE STANDARD DEVIATION ARE REPORTED, I.E., AVG (STD). THE BEST IS IN BOLD AND THE 2ND BEST IS UNDERLINED.

Feature Set	AUROC	Sensitivity	Specificity
XGBoost			
Subset 1	0.7061 (0.0265)	0.6093 (0.0382)	0.6898 (0.0289)
Subsets 1+2	0.7130 (0.0265)	0.6284 (0.0286)	0.6735 (0.0262)
Subsets 1+3	0.7398 (0.0297)	0.6688 (0.0399)	0.6920 (0.0233)
Subsets 1+2+3	0.7429 (0.0314)	0.6730 (0.0412)	0.6966 (0.0198)
Multi-Modal RNN			
Subset 1	0.7097 (0.0271)	0.6264 (0.0478)	0.6950 (0.0721)
Subsets 1+2	0.7100 (0.0301)	0.6371 (0.0519)	0.6745 (0.0757)
Subsets 1+3	0.7794 (0.0350)	0.7344 (0.0617)	0.6983 (0.0219)
Subsets 1+2+3	0.7716 (0.0227)	0.7048 (0.0352)	0.7106 (0.0304)

With the above observations, in the following, we use only feature subsets 1 and 3 to compare the proposal NPRL with other baselines, which is also fair to LSTM w/ Attention [21] that can only use temporal data. Table III shows the results. We can first observe through AUROC that most RNNs can outperform the traditional machine learning method XGBoost with a large margin as observed in the literature by capturing the temporal patterns. Except for RNN with the class-balanced loss from [24], which is because our problem is too imbalanced to be alleviated by the balanced loss. As well, balanced loss is often more helpful for multi-class classification. This is further demonstrated by additionally applying undersampling at first to bring the majority class to 2,600 instances, which improves the performance of the class-balanced RNN significantly. However, we can observe that the model is biased

² <https://github.com/ML4UWHealth/NPRL>

TABLE V

AVERAGE PERFORMANCE COMPARISON USING SUBSETS 1 AND 3. THE AVERAGE PERFORMANCE OVER 5 TEST FOLDS WITH THE STANDARD DEVIATION ARE REPORTED, I.E., AVG (STD). THE BEST IS IN BOLD AND THE 2ND BEST IS UNDERLINED.

Model	AUROC	Sensitivity	Specificity
XGBoost	0.7398 (0.0297)	0.6688 (0.0399)	0.6920 (0.0233)
LSTM w/ Attention [21]	0.7610 (0.0260)	0.8088 (0.0605)	0.5974 (0.0705)
RNN	0.7794 (0.0350)	0.7344 (0.0617)	0.6983 (0.0219)
RNN with Class Balanced Loss	0.7267 (0.0318)	0.6710 (0.0417)	0.6860 (0.0761)
RNN with Class Balanced Loss (Undersampling)	0.7733 (0.0293)	0.6900 (0.0910)	0.7075 (0.0707)
RNN with NPRL	0.7876 (0.0198)	0.8280 (0.0558)	0.6123 (0.0772)

to non-sepsis cases, where the performance on true negatives is significantly higher but the performance on true positives is quite low. On the contrary, [21] significantly improves the performance on true positives for sepsis instances, but sacrifices too much on non-sepsis instances, which results in lower AUROC.

Nevertheless, our proposal with nightly profile presentation learning provides the best AUROC and further improves the sensitivity of the model in predicting the target event significantly, without sacrificing the performance on non-sepsis cases that much. Considering that ICUs prefer to use machine learning models as a pre-screening assistant, they care more about missed positive cases (i.e., never being checked in-depth by clinicians). We can also observe the similar results based on the average performance of the 5 test folds reported in Table V. This further demonstrates the efficacy of our proposed method to address serious class imbalance problem for early sepsis prediction.

V. CONCLUSION

In this study, we propose a novel prediction setup for early sepsis onset anticipation by machine learning, which is more applicable and deploy-able as a pre-screening assistant in ICUs. To solve the problem of serious class imbalance in the dataset, we propose to do nightly patient profile representation learning that uniquely identify each patient at each night. Based on such pre-trained model, the diversity between different examples can be highly preserved, so as to alleviate the problem of major class domination. The proposed methodology demonstrates promising results in improving the overall AUROC performance and model sensitivity in early prediction of sepsis, which will assist ICU staff in intervening early. In future work, we will also need to deploy the model and evaluate its applicability in a live setting.

However, the high sensitivity comes from sacrificing of the specificity, although the highest AUROC is achieved and all trauma patients in the ICU will require attention during morning rounds, where a false positive will impose only a tiny amount of additional attention. One potential factor is the over and under sampling procedure adopted, where certain negative ones have been ignored. Therefore, a future work is to eliminate the over and under sampling, but still addressing the class imbalance problem using a new loss function inspired by our theoretically analysis from NPRL. Moreover, data augmentation on images is an essential technique for the

success of self-supervised learning in computer vision, which will be explored for EHR data as our future work for further improvement. Finally, we were not able to demonstrate our proposal on public data like MIMIC III [50] due to the lack of a precise sepsis onset timestamp, whose assignment is a valuable future work.

REFERENCES

- [1] Y. Sakr, U. Jaschinski, X. Wittebole, T. Szakmany, J. Lipman, S. A. Namendys Silva, I. Martin-Loeches, M. Leone, M.-N. Lupu, and J.-L. Vincent, "Sepsis in intensive care unit patients: Worldwide data from the intensive care over nations audit," *Open forum infectious diseases*, vol. 5, no. 12, pp. ofy313–ofy313, 2018.
- [2] J. C. Mira, J. Cuschieri, T. Ozrazgat-Baslanti, Z. Wang, G. L. Ghita, T. J. Loftus, J. A. Stortz, S. L. Raymond, J. D. Lanz, L. V. Hennessy *et al.*, "The epidemiology of chronic critical illness after severe traumatic injury at two level one trauma centers," *Critical care medicine*, vol. 45, no. 12, p. 1989, 2017.
- [3] J. A. Stortz, J. C. Mira, S. L. Raymond, T. J. Loftus, T. Ozrazgat-Baslanti, Z. Wang, G. L. Ghita, C. Leeuwenburgh, M. S. Segal, A. Bihorac *et al.*, "Benchmarking clinical outcomes and the immunocatabolic phenotype of chronic critical illness after sepsis in surgical intensive care unit patients," *The journal of trauma and acute care surgery*, vol. 84, no. 2, p. 342, 2018.
- [4] J. A. Stortz, T. J. Murphy, S. L. Raymond, J. C. Mira, R. Ungaro, M. L. Dirain, D. C. Nacionales, T. J. Loftus, Z. Wang, T. Ozrazgat-Baslanti *et al.*, "Evidence for persistent immune suppression in patients who develop chronic critical illness after sepsis," *Shock (Augusta, Ga.)*, vol. 49, no. 3, p. 249, 2018.
- [5] R. B. Hawkins, S. L. Raymond, J. A. Stortz, H. Horiguchi, S. C. Brakenridge, A. Gardner, P. A. Efron, A. Bihorac, M. Segal, F. A. Moore *et al.*, "Chronic critical illness and the persistent inflammation, immunosuppression, and catabolism syndrome," *Frontiers in Immunology*, vol. 9, p. 1511, 2018.
- [6] T. Wang, A. Derhovanessian, S. De Cruz, J. A. Belperio, J. C. Deng, and G. S. Hoo, "Subsequent infections in survivors of sepsis: epidemiology and outcomes," *Journal of intensive care medicine*, vol. 29, no. 2, pp. 87–95, 2014.
- [7] A. K. Gardner, G. L. Ghita, Z. Wang, T. Ozrazgat-Baslanti, S. L. Raymond, R. T. Mankowski, B. A. Brumback, P. A. Efron, A. Bihorac, F. A. Moore *et al.*, "The development of chronic critical illness determines physical function, quality of life, and long-term survival among early survivors of sepsis in surgical intensive care units," *Critical care medicine*, vol. 47, no. 4, p. 566, 2019.
- [8] S.-M. Ou, H. Chu, P.-W. Chao, Y.-J. Lee, S.-C. Kuo, T.-J. Chen, C.-M. Tseng, C.-J. Shih, and Y.-T. Chen, "Long-term mortality and major adverse cardiovascular events in sepsis survivors: a nationwide population-based study," *American journal of respiratory and critical care medicine*, vol. 194, no. 2, pp. 209–217, 2016.
- [9] R. Liu, J. L. Greenstein, S. J. Granite, J. C. Fackler, M. M. Bembea, S. V. Sarma, and R. L. Winslow, "Data-driven discovery of a novel sepsis pre-shock state predicts impending septic shock in the icu," *Scientific reports*, vol. 9, no. 1, pp. 1–9, 2019.
- [10] A. Kumar, D. Roberts, K. E. Wood, B. Light, J. E. Parrillo, S. Sharma, R. Suppes, D. Feinstein, S. Zanotti, L. Taiberg *et al.*, "Duration of hypotension before initiation of effective antimicrobial therapy is the critical determinant of survival in human septic shock," *Critical care medicine*, vol. 34, no. 6, pp. 1589–1596, 2006.

- [11] D. B. Page, J. P. Donnelly, and H. E. Wang, "Community-, healthcare- and hospital-acquired severe sepsis hospitalizations in the university healthsystem consortium," *Critical care medicine*, vol. 43, no. 9, p. 1945, 2015.
- [12] E. Eguia, A. N. Cobb, M. S. Baker, C. Joyce, E. Gilbert, R. Gonzalez, M. Afshar, and M. M. Churpek, "Risk factors for infection and evaluation of sepsis-3 in patients with trauma," *The American Journal of Surgery*, vol. 218, no. 5, pp. 851–857, 2019.
- [13] J. Eriksson, M. Eriksson, O. Brattström, E. Hellgren, O. Friman, A. Gidlöf, E. Larsson, and A. Oldner, "Comparison of the sepsis-2 and sepsis-3 definitions in severely injured trauma patients," *Journal of Critical Care*, vol. 54, pp. 125–129, 2019.
- [14] F. Khoshnevisan, *A Variational Recurrent Adversarial Multi-Source Domain Adaptation Framework for Septic Shock Early Prediction Across Medical Systems*. North Carolina State University, 2021.
- [15] J. Futoma, S. Hariharan, K. Heller, M. Sendak, N. Brajer, M. Clement, A. Bedoya, and C. O'Brien, "An improved multi-output gaussian process rnn with real-time validation for early sepsis detection," in *Machine Learning for Healthcare Conference*. PMLR, 2017, pp. 243–254.
- [16] A. Khojandi, V. Tansakul, X. Li, R. S. Koszalinski, and W. Paiva, "Prediction of sepsis and in-hospital mortality using electronic health records," *Methods Inf Med*, vol. 57, no. 04, pp. 185–193, 2018.
- [17] M. Scherpf, F. Gräßer, H. Malberg, and S. Zaunseder, "Predicting sepsis with a recurrent neural network using the mimic iii database," *Computers in biology and medicine*, vol. 113, p. 103395, 2019.
- [18] S. M. Lauritsen, M. Kristensen, M. V. Olsen, M. S. Larsen, K. M. Lauritsen, M. J. Jørgensen, J. Lange, and B. Thiesson, "Explainable artificial intelligence model to predict acute critical illness from electronic health records," *Nature communications*, vol. 11, no. 1, pp. 1–11, 2020.
- [19] S. Ramchand, "An ensemble lstm for rare event detection," 2020.
- [20] M. Moor, M. Horn, B. Rieck, D. Roqueiro, and K. Borgwardt, "Temporal convolutional networks and dynamic time warping can drastically improve the early prediction of sepsis," *arXiv preprint arXiv:1902.01659*, 2019.
- [21] D. A. Kaji, J. R. Zech, J. S. Kim, S. K. Cho, N. S. Dangayach, A. B. Costa, and E. K. Oermann, "An attention based deep learning model of clinical events in the intensive care unit," *PloS one*, vol. 14, no. 2, p. e0211057, 2019.
- [22] S. M. Lauritsen, M. E. Kalør, E. L. Kongsgaard, K. M. Lauritsen, M. J. Jørgensen, J. Lange, and B. Thiesson, "Early detection of sepsis utilizing deep learning on electronic health record event sequences," *Artificial Intelligence in Medicine*, vol. 104, p. 101820, 2020.
- [23] J. M. Johnson and T. M. Khoshgoftaar, "Survey on deep learning with class imbalance," *Journal of Big Data*, vol. 6, no. 1, pp. 1–54, 2019.
- [24] Y. Cui, M. Jia, T.-Y. Lin, Y. Song, and S. Belongie, "Class-balanced loss based on effective number of samples," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 9268–9277.
- [25] H. Liu, J. Z. HaoChen, A. Gaidon, and T. Ma, "Self-supervised learning is more robust to dataset imbalance," in *International Conference on Learning Representations*, 2022. [Online]. Available: <https://openreview.net/forum?id=4AZz9osqar>
- [26] S. M. Lauritsen, B. Thiesson, M. J. Jørgensen, A. H. Riis, U. S. Espelund, J. B. Weile, and J. Lange, "The consequences of the framing of machine learning risk prediction models: evaluation of sepsis in general wards," *arXiv preprint arXiv:2101.10790*, 2021.
- [27] J. Futoma, S. Hariharan, and K. Heller, "Learning to detect sepsis with a multitask gaussian process rnn classifier," in *International Conference on Machine Learning*. PMLR, 2017, pp. 1174–1182.
- [28] H. J. Kam and H. Y. Kim, "Learning representations for the early detection of sepsis with deep neural networks," *Computers in biology and medicine*, vol. 89, pp. 248–255, 2017.
- [29] S. Nemati, A. Holder, F. Razmi, M. D. Stanley, G. D. Clifford, and T. G. Buchman, "An interpretable machine learning model for accurate prediction of sepsis in the icu," *Critical care medicine*, vol. 46, no. 4, p. 547, 2018.
- [30] F. Van Wyk, A. Khojandi, and R. Kamaleswaran, "Improving prediction performance using hierarchical analysis of real-time data: a sepsis case study," *IEEE journal of biomedical and health informatics*, vol. 23, no. 3, pp. 978–986, 2019.
- [31] J. Guillén, J. Liu, M. Furr, T. Wang, S. Strong, C. C. Moore, A. Flower, and L. E. Barnes, "Predictive models for severe sepsis in adult icu patients," in *2015 Systems and Information Engineering Design Symposium*. IEEE, 2015, pp. 182–187.
- [32] C. Barton, U. Chettipally, Y. Zhou, Z. Jiang, A. Lynn-Palevsky, S. Le, J. Calvert, and R. Das, "Evaluation of a machine learning algorithm for up to 48-hour advance prediction of sepsis using six vital signs," *Computers in biology and medicine*, vol. 109, pp. 79–84, 2019.
- [33] A. Wong, E. Otles, J. P. Donnelly, A. Krumm, J. McCullough, O. DeTroyer-Cooley, J. Pestreue, M. Phillips, J. Konye, C. Penzoza et al., "External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients," *JAMA Internal Medicine*, vol. 181, no. 8, pp. 1065–1070, 2021.
- [34] A. Teredesai, S. Huang, T. Stewart, J. Hu, A. Thakker, K. Stern, and G. E. O'Keefe, "Sub-sequence graph representation learning on high variability data for dynamic risk prediction in critical care," in *2022 IEEE International Conference on Big Data (Big Data)*. IEEE, 2022, pp. 2082–2092.
- [35] K. Ewig, X. Lin, T. Stewart, K. Stern, G. O'Keefe, A. Teredesai, and J. Hu, "Multi-subset approach to early sepsis prediction," *arXiv preprint arXiv:2304.06384*, 2023.
- [36] J. M. Johnson and T. M. Khoshgoftaar, "Survey on deep learning with class imbalance," *Journal of Big Data*, vol. 6, no. 1, pp. 1–54, 2019.
- [37] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [38] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using rnn encoder-decoder for statistical machine translation," *arXiv preprint arXiv:1406.1078*, 2014.
- [39] E. Choi, A. Schuetz, W. F. Stewart, and J. Sun, "Using recurrent neural network models for early detection of heart failure onset," *Journal of the American Medical Informatics Association*, vol. 24, no. 2, pp. 361–370, 2017.
- [40] F. Ma, R. Chitta, J. Zhou, Q. You, T. Sun, and J. 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*, 2017, pp. 1903–1911.
- [41] H. H. Sohrab, *Basic real analysis*. Springer, 2003, vol. 231.
- [42] A. Sinha, H. Namkoong, and J. C. Duchi, "Certifying some distributional robustness with principled adversarial training," in *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. [Online]. Available: <https://openreview.net/forum?id=Hk6kPgZA->
- [43] K. Stern, Q. Qiu, M. Weykamp, G. O'Keefe, and S. C. Brakenridge, "Defining posttraumatic sepsis for population-level research," *JAMA Network Open*, vol. 6, no. 1, pp. e2251445–e2251445, 2023.
- [44] M. Singer, C. S. Deutschman, C. W. Seymour, M. Shankar-Hari, D. Annane, M. Bauer, R. Bellomo, G. R. Bernard, J.-D. Chiche, C. M. Coopersmith et al., "The third international consensus definitions for sepsis and septic shock (sepsis-3)," *Jama*, vol. 315, no. 8, pp. 801–810, 2016.
- [45] A. F. Monegro, V. Muppidi, and H. Regunath, "Hospital acquired infections," in *StatPearls [Internet]*. StatPearls Publishing, 2020.
- [46] J. L. Vincent, R. Moreno, J. Takala, S. Willatts, A. De Mendonça, H. Bruining, C. K. Reinhart, P. M. Suter, and L. G. Thijs, "The sofa (sepsis-related organ failure assessment) score to describe organ dysfunction/failure. on behalf of the working group on sepsis-related problems of the european society of intensive care medicine," *Intensive care medicine*, vol. 22, no. 7, pp. 707–710, 1996.
- [47] H. D. Sesso, M. J. Stampfer, B. Rosner, C. H. Hennekens, J. M. Gaziano, J. E. Manson, and R. J. Glynn, "Systolic and diastolic blood pressure, pulse pressure, and mean arterial pressure as predictors of cardiovascular disease risk in men," *Hypertension*, vol. 36, no. 5, pp. 801–807, 2000.
- [48] S. A. Magder, "The highs and lows of blood pressure: toward meaningful clinical targets in patients with shock," *Critical care medicine*, vol. 42, no. 5, pp. 1241–1251, 2014.
- [49] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," in *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 2016, pp. 785–794.
- [50] E. J. Tsiklidis, T. Sinno, and S. L. Diamond, "Predicting risk for trauma patients using static and dynamic information from the mimic iii database," *Plos one*, vol. 17, no. 1, p. e0262523, 2022.