

---

# Benchmarking machine learning models on eICU critical care dataset

---

**Syedmostafa Sheikhalishahi, Vevake Balaraman**  
Fondazione Bruno Kessler Research Institute and University of Trento  
Trento, Italy  
{ssheikhalishahi, balaraman}@fbk.eu

**Venet Osmani**  
Fondazione Bruno Kessler Research Institute  
Trento, Italy  
vosmani@fbk.eu

## Abstract

Progress of machine learning in critical care has been difficult to track, in part due to absence of public benchmarks. Other fields of research (such as vision and NLP) have already established various competitions and benchmarks, whereas only recent availability of large clinical datasets has enabled the possibility of public benchmarks. Taking advantage of this opportunity, we propose a public benchmark suite to address four areas of critical care, namely mortality prediction, estimation of length of stay, patient phenotyping and risk of decompensation. We define each task and compare the performance of both clinical models as well as baseline and deep models using eICU critical care dataset of around 73,000 patients. Furthermore, we investigate the impact of numerical variables as well as handling of categorical variables for each of the defined tasks.

## 1 Introduction

Increasing availability of clinical data and advances in machine learning have addressed a wide range of healthcare problems, such as risk assessment and prediction both in acute, chronic and critical care. Critical care is an especially data-intensive field, as continuous monitoring of patients in Intensive Care Units (ICU) generates large streams of data that can then be harnessed by machine learning algorithms. However, progress in harnessing digital health data faces several obstacles, including reproducibility and comparison of results between competing models. While other areas of machine learning research, such as image and natural language processing have established a number of competitions and benchmarks (such as ILSVRC on N2C2, respectively), progress in machine learning for critical care has been difficult to measure, in part due to absence of public benchmarks. However, availability of large clinical data sets, including MIMIC III (Johnson et al., 2016) and more recently eICU (Pollard et al., 2018) are opening the possibility of establishing public benchmarks and consequently tracking the progress of machine learning models in critical care. In this paper, we propose a public benchmark suite to address four areas of critical care, namely mortality prediction, estimation of length of stay, patient phenotyping and risk of decompensation. We define each task and evaluate our algorithms on a dataset of 73,718 patients (containing 4,564,844 clinical records). While, there has been an initial work in this area that has focused on MIMIC III clinical dataset (Harutyunyan et al., 2019), our work is the first to focus on a multi-center intensive care unit dataset, the eICU database (Pollard et al., 2018).

The main contributions of this work include: i) we provide the baseline performance and compare it with our benchmark result, using a model based on bidirectional LSTM; ii) investigate impact of

	Overall	Dead at Hospital	Alive at Hospital
Admissions	73718	6167	67551
Age	62.41 [52-75]	68.12 [59-80]	61.8 [52-75]
Gender (F)	33544 (45.5)	2830 (45.8)	30714 (45.4)
<i>Ethnicity</i>			
Caucasian	56973 (77.2)	4866 (78.9)	52107 (77.1)
African American	7982 (10.8)	582 (9.4)	7400 (10.9)
Hispanic	2937 (3.98)	226 (3.6)	2711 (4)
Asian	1174 (1.59)	97 (1.5)	1077 (1.5)
Native American	413 (0.56)	42 (0.68)	371 (0.54)
Unknown	4239 (5.7)	354 (5.7)	3885 (5.7)
<i>Outcomes</i>			
Hospital LoS* (days)	5.29 [2.53-6.84]	3.9 [1.42-5.22]	5.41 [2.65-6.92]
ICU LoS* (days)	2.32 [1.01-2.91]	3.17 [1.19-4.43]	2.24 [1-2.83]
Hospital Death	6167 (8.36)	6167 (100)	-
ICU Death	4575 (6.2)	4575 (74.1)	-

Table 1: Characteristics and mortality outcome measures. \*LoS (Length of Stay). Continuous variables are presented as Median [Interquartile Range Q1–Q3]; binary or categorical variables as Count (%)

categorical and numerical variables on all four benchmarking tasks; iii) evaluate entity embedding for categorical variables, versus one hot encoding; iv) show that for some tasks the number of variables can be reduced significantly without greatly impacting prediction performance; and v) we report six evaluation metrics for each of the tasks, facilitating direct comparison with future results. The source code for our experiments will be made public at <https://github.com/eICU-benchmark> so that anyone with access to eICU database can replicate our experiments and build upon our work.

## 2 eICU dataset description and cohort selection

The eICU Collaborative Research Database (Pollard et al., 2018) is a multi-center intensive care unit database with high granularity data for over 200,000 admissions to ICUs monitored by eICU programs across The United States. The eICU database comprises 200,859 patient unit encounters for 139,367 unique patients admitted between 2014 and 2015 to hospitals located throughout the US. We selected adult patients only (age > 18) that had an ICU admission with at least 15 records, leading to 73,718 unique patients with median age of 62.41 years (Q1–Q3: 52-75), 45.5% female. Hospital mortality rate was 8.3% and average length of stay in hospital and in unit were 5.29 days and 3.9 days respectively (further details provided in Table 1). The final patient cohort contained 4,564,844 clinical records where we group these records on 1 hour window, impute the missing values based on the mean of that window and take the last valid record.

Out of 31 tables in the eICU (v1.0) database we selected variables from the following tables: *patient* (administrative information and patient demographics), *lab* (Laboratory measurements collected during routine care), *nurse charting* (bedside documentation) and *diagnosis* based on advice from a clinician as well as consistency with other similar tasks reported in Section 4. Selected variables are shown in Table 2.

## 3 Benchmarking experiments

### 3.1 Description of tasks

In this section, we define four different benchmark tasks, namely in-hospital mortality prediction, remaining length of stay forecasting, patient phenotyping, and risk of physiologic decompensation. After applying selection criteria, the resulting patient cohorts are outlined in Table 3

#### 3.1.1 Mortality prediction

Task	No. of patients	Clinical records
In-hospital Mortality	30,680	1,164,966
RLoS	73,389	3,054,314
Phenotyping	49,768	2,192,497
Physiologic Decompensation	55,933	2,800,711

Table 3: Number of patients and records in four tasks

In-hospital mortality is defined as the patient’s outcome at the hospital discharge. This is a binary classification task, where each data sample spans a 1-hour window. The cohort for this task was selected based on the presence of hospital discharge status in patients’ record and length of stay of at least 48 hours (we focus on prediction during the first 24 and 48 hours). This selection criteria resulted in 30,680 patients containing 1,164,966 records.

### 3.1.2 Length of stay prediction

Length of stay is one of the most important factors accounting for the overall hospital costs, as such its forecast could play an important role in healthcare management (Kılıç et al., 2019). Length of stay is estimated through analysis of events occurring within a fixed time-window, once every hour from the initial ICU admission. This is a regression task, where we use 20 clinical variables described in Table 2. For this cohort we selected patients whose length of stay was present in their records with a duration of more than 24 hours. These selection criteria resulted in 73,389 ICU stays, containing 3,054,314 records. The mean length of stay was 1.86 days with standard deviation of 1.94 days, as shown in Table 1.

### 3.1.3 Phenotyping

Phenotyping refers to the prediction for the diseases (ICD-9 codes). Since any given patient may have more than one ICD-9 code, this is defined as a multi-label classification problem. The dataset contains 767 unique ICD codes, which are grouped into 25 categories shown in Table 6. The cohort for this task, considering initial inclusion criteria as well as recorded diagnosis during the ICU stay, results in 49,768 patients.

### 3.1.4 Physiologic Decompensation

There are a number of ways to define decompensation, however in clinical setting majority of early warning systems, such as National Early Warning Score (NEWS) (McGinley and Pearse, 2012) are based on prediction of mortality within the next time window (such as 24 hours after the assessment). Following suit and keeping consistent with previously published benchmarks (Harutyunyan et al., 2019), we also define decompensation as a binary classification problem, where the target label indicates whether the patient dies within the next 24 hours. The cohort for this task results in 55,933 patients (2,800,711 records), where decompensation rate is around 6.5% (3664 patients).

Variable	Data Type
Heart rate	Numerical
Mean arterial pressure	Numerical
Diastolic blood pressure	Numerical
Systolic blood pressure	Numerical
O2	Numerical
Respiratory rate	Numerical
Temperature	Numerical
Glucose	Numerical
FiO2	Numerical
pH	Numerical
Height	Numerical
Weight	Numerical
Age	Numerical
Admission diagnosis	Categorical
Ethnicity	Categorical
Gender	Categorical
Glasgow Coma Score Total	Categorical
Glasgow Coma Score Eyes	Categorical
Glasgow Coma Score Motor	Categorical
Glasgow Coma Score Verbal	Categorical

Table 2: Selected variables for all the four tasks

## 3.2 Prediction algorithms

### 3.2.1 Baselines

We compare our model with two standard baseline approaches namely, logistic regression (LR) and a 1-layer artificial neural network (ANN). Both these approaches use a flattened representation of all the features concatenated in the order of the timestep. The embeddings for these models are learned in the same way as for the proposed BiLSTM model as explained in Section 3.2.2.

### 3.2.2 Deep Learning models

In this section, we describe the selected clinical variables, approaches to represent these variables as well as baseline and deep models used in this study. The architecture of this work consists of three modules, namely input module, encoder module and output module as shown in Figure 1.

**Input representation:** We process and model both numerical and categorical variables separately, as shown in Table 2. Categorical variables are represented using either one-hot encoding (OHE) or entity embedding (EE). OHE is the baseline approach that converts the variables into binary representation. Using this approach for our 7 categorical variables results in 429 unique records, rendering a large sparse matrix. In response, we represent each variable as an embedding and compare the performance with the OHE approach. We use entity embedding (Guo and Berkahn, 2016), where each categorical variable in the dataset is mapped to a vector and the corresponding embedding is added to the patient’s record. This entity embedding is learned by the neural network during the training phase along with other parameters. So the final representation of the input at time  $t$  is as follows:

$$x_t = [Num_t; U(Cat_t)]$$

where  $Num_t$  is the numerical variable,  $Cat_t$  is the categorical variable at time  $t$  and  $U$  is the embedding matrix learned by the model.

**Encoder:** To capture sequential dependency in our data, we use Recurrent Neural Network (RNN) that resemble a chain of repeating modules to efficiently model sequential data (Rumelhart et al., 1988). They take sequential data  $X = (x_1, x_2, \dots, x_n)$  as input and provide a hidden representation  $H = (h_1, h_2, \dots, h_n)$  which captures the information at every time step in the input. Formally,

$$h_t = f(x_t + Wh_{t-1})$$

where  $x_t$  is the input at time  $t$ ,  $W$  is the parameter of RNN learned during training and  $f$  is a non-linear operation such as sigmoid, tanh or ReLU.

A drawback of regular RNNs is that the input sequence is fed in one direction, normally from past to future. In order to capture both past and future context, we use a Bidirectional Long Short Term Memory (BiLSTM) (Schuster and Paliwal, 1997) (Hochreiter and Schmidhuber, 1997) for our model, which processes the input in both forward and backward direction. Using a BiLSTM the model is able to capture the context of a record not only by its preceding records but also with the following records, allowing the model to produce more informed predictions. The input at time  $t$  is represented by both its forward context  $\vec{h}_t$  and backward context  $\overleftarrow{h}_t$  as  $h_t = [\vec{h}_t; \overleftarrow{h}_t]$ . Similarly, the representation of the completed patient record is given by  $h_T = [\vec{h}_n; \overleftarrow{h}_1]$ .

**Output:** The choice of output layer is based on whether the benchmarking task is a regression or a classification task.

Remaining length of stay (RLoS) prediction is a regression task, in which we predict the RLoS record-wise. That is, each patient record is fed to the model to predict RLoS for that specific time step. This task is realized using a many to many architecture, where we assign a label to each patient record. The score for this task is obtained using:

$$\hat{y}_t = ReLU(W \cdot h_t) \tag{1}$$

where  $y_t$  is the RLoS predicted and  $ReLU$  is the non-linear activation function used as the prediction of RLoS cannot be negative.

In-hospital mortality and decompensation are binary classification tasks. For the in-hospital mortality the many to one architecture is applied and the classifier is as follows:

$$\hat{y} = \sigma(W \cdot h_T) \quad (2)$$

For the decompensation task, a many to many architecture is applied. Prediction at each-time step is treated as a binary classification and the classifier is defined as:

$$\hat{y}_t = \sigma(W \cdot h_t) \quad (3)$$

Phenotyping is defined as a multi-label task with 25 binary classifiers for each phenotype, and the score for the task is obtained using:

$$\hat{y}_t^n = \sigma(W_n \cdot h_t) \quad (4)$$

where  $t$  is the time step and  $n$  is the phenotype being predicted and  $W_n$  is the model parameter.

### 3.3 Results

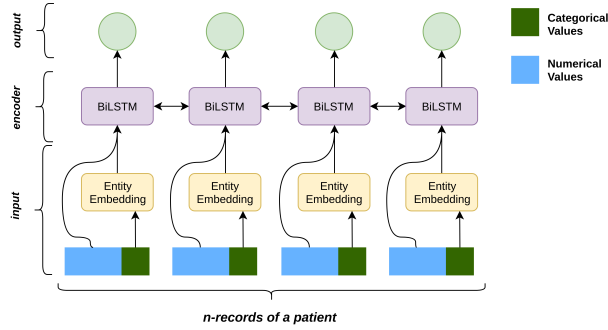


Figure 1: Model architecture.

In this section, we report benchmarking results of methods and prediction algorithms, focusing on answering the following questions: (a) How does performance of different models compare to performance of clinical scoring systems?; and (b) What is impact on prediction performance when using different feature sets, such as categorical and numerical variables, solely categorical and solely numerical variables?

We evaluate our model through a random 80/20 train and test split using the following evaluation metrics: for the regression tasks we report coefficient of determination

$R^2$ , and Mean Absolute Error (MAE), while for the classification tasks we report AUROC, AUPRC, Specificity, Sensitivity, Positive Predictive Value (PPV) and Negative Predictive Value (NPV).

#### 3.3.1 Mortality prediction

Results from this task indicate that the proposed approach of learning embeddings for categorical variables is more effective than OHE representation. This holds true for both baseline models (LR and ANN) as well as BiLSTM model, reflected in the prediction performance of each model. Furthermore, BiLSTM based model outperforms all the other approaches in predicting mortality in both 24 hour window and 48 hour window as shown in Table 4. It is interesting to note that using only categorical variables (reducing the number of variables from 20 to only 7) with embedding provides a better performance than using numerical variables only (AUROC 79.3 vs. 74 and 80.1 vs 76.7 - for the first 24h and 48h, respectively). This holds true also for categorical variables with one hot encoding. These results suggest that entity embedding of categorical features in vector space is more effective in the prediction of mortality.

#### 3.3.2 Length of stay prediction

Predicting Length of Stay (LoS) requires capturing temporal dependencies between each time-step. For this reason baseline models perform poorly due the lack of explicit modelling of temporal dependencies. The proposed BiLSTM model on the other hand is able to capture this dependency effectively and outperforms the baseline models as shown in Table 5. We can also see that the numerical variables are the most effective in prediction of LoS. However, using categorical variables encoded with OHE reduces the model performance, while entity embedding improves MAE, Kappa and AUROC measures.

Data	Model	Num.	Cat.	Repn.	AUROC	AUPRC	Spec.	Sens.	PPV	NPV
First 24 hours	APACHE	✓	✓	Not spec.	66.1	53	97	35	63	91
	LR	✓	✓	EMB	80.8	46.7	96	32	56	91
	ANN	✓	✓	EMB	82.9	50.9	97	34	62	91
	BiLSTM	✓	×	×	74	35.6	80	54	29	92
	BiLSTM	×	✓	OHE	77.7	42.6	78	61	29	93
	BiLSTM	×	✓	EMB	79.3	44.5	90	45	39	92
	BiLSTM	✓	✓	OHE	75.4	36.6	62	75	22	94
	<b>BiLSTM</b>	✓	✓	EMB	<b>83.6</b>	50.1	90	55	44	93
First 48 hours	LR	✓	✓	EMB	85.4	55.07	97	36	65	91
	ANN	✓	✓	EMB	84.8	54.25	95	47	60	93
	BiLSTM	✓	×	×	76.7	39.80	97	25	52	90
	BiLSTM	×	✓	OHE	78.9	46.8	85	56	35	93
	BiLSTM	×	✓	EMB	80.1	47.32	79	65	31	94
	BiLSTM	✓	✓	OHE	71.3	34.12	61	67	20	93
	BiLSTM	✓	✓	EMB	<b>86</b>	54.45	84	71	40	95

Table 4: In-hospital mortality prediction during first 24 and 48 hours in ICU. (*Num.* and *Cat.* indicate presence of numerical and categorical variables respectively. *Repn.* indicates representation of categorical variables, either One Hot Encoding (OHE) or embedding (EMB) )

Data	Model	Num.	Cat.	Repn.	$R^2$	MAE [Day]
In ICU unit	LR	✓	✓	EMB	0.03	1.25
	ANN	✓	✓	EMB	0.032	1.253
	BiLSTM	✓	×	×	0.79	0.47
	BiLSTM	×	✓	EMB	0.75	0.50
	BiLSTM	✓	✓	OHE	0.70	0.52
	<b>BiLSTM</b>	✓	✓	EMB	0.74	<b>0.46</b>

Table 5: Length of stay in hospital prediction, evaluated using Mean Absolute Error (MAE)

### 3.3.3 Phenotyping prediction

For the phenotyping task, we focus on comparing performance (AUROC) of the proposed model on different subset of features, namely numerical versus categorical variables. Using only the categorical features, modelled as entity embeddings shows a significantly higher performance (0.84) compared to using only the numerical features (0.56) as outlined in Figure 6. Clearly categorical features are more effective in representing patients’ phenotype, since integrating both of the subsets does not significantly improve the result (0.86 from 0.84). In this task there is a wide difference between performance of the model on individual diseases, varying from 0.61 (diabetes mellitus without complications) to 0.96 (acute cerebrovascular disease). As a general trend prediction performance on acute diseases is higher (0.83) than that on chronic diseases (0.72). This may be due to slow-progressing nature of chronic diseases, where recorded ICU data is relatively short and thus unable to fully capture events related to chronic diseases.

### 3.3.4 Decompensation prediction

As mentioned in Section 3.1.4, decompensation is related to mortality prediction with the difference that we predict whether the patients survives in the next 24 hours, given the current time step. As such, time-dependence is critical. Since 3 categorical variables (out of 7) are time-independent and only 4 are time-dependent, they pose a difficult challenge for the model to be able to predict decompensation using only the time-independent categorical variables. For this reason, the model with only numerical variables outperforms other approaches as shown in Table 7.

Phenotype	Prevalence	Type	Num & cat	Num.	Cat.
Respiratory failure; insufficiency; arrest	0.241	acute	0.86	0.61	0.85
Fluid and electrolyte disorders	0.155	acute	0.70	0.57	0.70
Septicemia	0.144	acute	0.90	0.65	0.89
Acute and unspecified renal failure	0.141	acute	0.75	0.55	0.73
Pneumonia	0.120	acute	0.87	0.62	0.87
Acute cerebrovascular disease	0.108	acute	<b>0.96</b>	<b>0.66</b>	<b>0.95</b>
Acute myocardial infarction	0.089	acute	0.92	0.58	0.91
Gastrointestinal hemorrhage	0.079	acute	0.92	0.53	0.91
Shock	0.068	acute	0.85	0.62	0.82
Pleurisy; pneumothorax; pulmonary collapse	0.039	acute	0.62	0.50	0.63
Other lower respiratory disease	0.029	acute	0.85	0.51	0.85
Complications of surgical	0.011	acute	0.71	0.54	0.73
Other upper respiratory disease	0.007	acute	0.91	0.55	0.92
<i>Macro-average (acute diseases)</i>	-	-	0.83	-	-
Hypertension with complications	0.018	chronic	0.85	0.47	0.83
Essential hypertension	0.203	chronic	0.68	0.62	0.64
Chronic kidney disease	0.103	chronic	0.67	0.49	0.62
Chronic obstructive pulmonary disease	0.093	chronic	0.77	0.54	0.77
Disorders of lipid metabolism	0.054	chronic	0.72	0.54	0.71
Coronary atherosclerosis and related	0.041	chronic	0.79	0.56	0.79
Diabetes mellitus without complication	0.006	chronic	0.61	0.52	0.51
<i>Macro-average (chronic diseases)</i>	-	-	0.72	-	-
Cardiac dysrhythmias	0.165	mixed	0.72	0.60	0.69
Congestive heart failure; non hypertensive	0.105	mixed	0.79	0.62	0.78
Diabetes mellitus with complications	0.046	mixed	0.89	0.75	0.85
Other liver diseases	0.038	mixed	0.80	0.55	0.78
Conduction disorders	0.012	mixed	0.81	0.58	0.82
<i>Macro-average (mixed diseases)</i>	-	-	0.80	-	-
<i>Micro-average (all diseases)</i>	-	-	0.86	0.56	0.84
<i>Macro-average (all diseases)</i>	-	-	0.80	0.57	0.78

Table 6: Phenotyping task on eICU (reported scores are AUROC)

Data	Model	Num.	Cat.	Repn.	AUROC	AUPRC	Spec.	Sens.	PPV	NPV
In ICU unit	LR	✓	✓	EMB	57	64.9	6	99	62	84
	ANN	✓	✓	EMB	66.4	71.8	22	94	65	71
	BiLSTM	✗	✓	EMB	49.3	38	100	0	0	65
	BiLSTM	✓	✓	EMB	97.4	94.4	92	91	86	95
	BiLSTM	✓	✗	✗	<b>98</b>	95.6	94	90	88	95

Table 7: Decompensation risk prediction in eICU

## 4 Related work

A number of scoring systems have been developed for mortality prediction, including Acute Physiology and Chronic Health Evaluation (APACHE III (Knaus et al., 1991), APACHE IV (Zimmerman et al., 2006)) and Simplified Acute Physiology Score (Le Gall et al., 1993) (SAPS II, SAPS III). Most of these scoring systems use logistic regression to identify predictive features to establish these scoring systems.

Providing an accurate prediction of mortality risk for patients admitted to ICU using the first 24/48 hours of ICU data could make the decision making easier and reduce the healthcare costs. In this regard, recent advanced techniques in artificial intelligence showed to outperform the conventional machine learning and clinical prediction techniques such as APACHE and SAPS (Harutyunyan et al., 2019) (Purushotham et al., 2018) (Lipton et al., 2015). Mortality prediction has been a popular application for deep learning researchers in recent years, though model architecture and problem

definition vary widely. Convolutional neural network and gradient boosted tree algorithm have been used by Darabi et al. (Darabi et al., 2018), in order to predict long-term mortality risk (30 days) on a subset of MIMIC-III dataset. Similarly, Celi et al. (Celi et al., 2012) developed mortality prediction models based on a subset of MIMIC database using logistic regression, Bayesian network and artificial neural network.

Harutyunyan et al. (Harutyunyan et al., 2019) developed a deep learning model based on RNN LSTM called multi-task RNN, in order to predict mortality prediction in hospital, decompensation, phenotyping and length of stay in ICU unit. The proposed model was applied on MIMIC-III dataset. Similarly, Purushotham et al (Purushotham et al., 2018) have done a comprehensive benchmark of several machine learning and deep learning models trained on MIMIC-III for various tasks, while results showing deep models consistently outperforming the conventional machine learning models and the scoring systems.

Previous work (Purushotham et al., 2018)(Harutyunyan et al., 2019) has shown that deep learning models obtain good results on forecasting length of stay in ICU. In this regards, Tu et al (Tu and Guerriere, 1993) applied neural network based methods on a Canadian private dataset which includes patients with cardiac surgery. The proposed model was able to detect the patient with low, intermediate and high prolonged stay in ICU. Deep learning methods have been applied to predict phenotyping by Razavian et al (Razavian et al., 2016) and Lipton et al (Lipton et al., 2015). While the first trained LSTM and CNN for prediction of 133 diseases based on 18 laboratory tests on a private dataset including 298k patients, the latter applied an RNN LSTM on a private pediatric intensive care unit (PICU) dataset in order to classify 128 diagnoses given 13 clinical measurements.

In terms of physiologic decompensation Xu et al (Xu et al., 2018), proposed an attention based model which outperformed several machine learning models in order to predict the decompensation event. They evaluated their proposed model on MIMIC-III Waveform Database Matched Subset. Similarly, (Harutyunyan et al., 2019) proposed decompensation prediction as one of the tasks in their multi-task benchmark.

## 5 Conclusion

In this study we have described four standardised benchmarks in critical care research. Our definition of benchmark tasks is consistent with previously published benchmarks, however we focus on the more recent eICU database, where clinical data has been collected from multiple ICU centres across the United States that may result in lower systematic bias. We provided a set of baselines for our benchmarks and show that bi-directional LSTM significantly outperforms linear models, especially in tasks with temporal dependencies, such as length of stay. Of note is the impact of entity embedding of categorical variables in further improving the performance of our LSTM-based model. We believe that our work will provide a solid basis to further improve critical care decision making and we provide the source code for other researchers that wish to replicate our experiments and build upon our results.

## Acknowledgements

We gratefully acknowledge clinical input provided by Monica Moz, MD in both cohort selection as well as variable ranking and selection.

## References

- Celi, L. A., Galvin, S., Davidzon, G., Lee, J., Scott, D., and Mark, R. (2012). A database-driven decision support system: customized mortality prediction. *Journal of personalized medicine*, 2(4):138–148.
- Darabi, H. R., Tsinis, D., Zecchini, K., Whitcomb, W. F., and Liss, A. (2018). Forecasting mortality risk for patients admitted to intensive care units using machine learning. *Procedia Computer Science*, 140:306 – 313. Cyber Physical Systems and Deep Learning Chicago, Illinois November 5-7, 2018.



- Guo, C. and Berkhahn, F. (2016). Entity embeddings of categorical variables. *arXiv preprint arXiv:1604.06737*.
- Harutyunyan, H., Khachatryan, H., Kale, D. C., Ver Steeg, G., and Galstyan, A. (2019). Multitask learning and benchmarking with clinical time series data. *Scientific data*, 6(1):96.
- Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Johnson, A. E., Pollard, T. J., Shen, L., Li-wei, H. L., Feng, M., Ghassemi, M., Moody, B., Szolovits, P., Celi, L. A., and Mark, R. G. (2016). MIMIC-III, a freely accessible critical care database. *Scientific data*, 3:160035.
- Kılıç, M., Yüzkat, N., Soyaloğlu, C., and Gülhaş, N. (2019). Cost analysis on intensive care unit costs based on the length of stay. *Turkish journal of anaesthesiology and reanimation*, 47(2):142.
- Knaus, W. A., Wagner, D. P., Draper, E. A., Zimmerman, J. E., Bergner, M., Bastos, P. G., Sirio, C. A., Murphy, D. J., Lotring, T., Damiano, A., et al. (1991). The apache iii prognostic system: risk prediction of hospital mortality for critically ill hospitalized adults. *Chest*, 100(6):1619–1636.
- Le Gall, J.-R., Lemeshow, S., and Saulnier, F. (1993). A new simplified acute physiology score (saps ii) based on a european/north american multicenter study. *Jama*, 270(24):2957–2963.
- Lipton, Z. C., Kale, D. C., Elkan, C., and Wetzell, R. (2015). Learning to diagnose with lstm recurrent neural networks. *arXiv preprint arXiv:1511.03677*.
- McGinley, A. and Pearse, R. M. (2012). A national early warning score for acutely ill patients.
- Pollard, T. J., Johnson, A. E., Raffa, J. D., Celi, L. A., Mark, R. G., and Badawi, O. (2018). The eicu collaborative research database, a freely available multi-center database for critical care research. *Scientific data*, 5.
- Purushotham, S., Meng, C., Che, Z., and Liu, Y. (2018). Benchmarking deep learning models on large healthcare datasets. *Journal of Biomedical Informatics*, 83(April):112–134.
- Razavian, N., Marcus, J., and Sontag, D. (2016). Multi-task prediction of disease onsets from longitudinal laboratory tests. In *Machine Learning for Healthcare Conference*, pages 73–100.
- Rumelhart, D. E., Hinton, G. E., Williams, R. J., et al. (1988). Learning representations by back-propagating errors. *Cognitive modeling*, 5(3):1.
- Schuster, M. and Paliwal, K. K. (1997). Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing*, 45(11):2673–2681.
- Tu, J. V. and Guerriere, M. R. (1993). Use of a neural network as a predictive instrument for length of stay in the intensive care unit following cardiac surgery. *Computers and Biomedical Research*, 26(3):220 – 229.
- Xu, Y., Biswal, S., Deshpande, S. R., Maher, K. O., and Sun, J. (2018). Raim: Recurrent attentive and intensive model of multimodal patient monitoring data. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2565–2573. ACM.
- Zimmerman, J. E., Kramer, A. A., McNair, D. S., and Malila, F. M. (2006). Acute physiology and chronic health evaluation (apache) iv: hospital mortality assessment for today’s critically ill patients. *Critical care medicine*, 34(5):1297–1310.