

RESEARCH ARTICLE

MGP-AttTCN: An interpretable machine learning model for the prediction of sepsis

Margherita Rosnati¹, Vincent Fortuin^{2*}¹ Department of Computing, Imperial College London, London, United Kingdom, ² Department of Computer Science, ETH Zürich, Zürich, Switzerland* fortuin@inf.ethz.ch

OPEN ACCESS

Citation: Rosnati M, Fortuin V (2021) MGP-AttTCN: An interpretable machine learning model for the prediction of sepsis. PLoS ONE 16(5): e0251248. <https://doi.org/10.1371/journal.pone.0251248>

Editor: Ivan Olier, Liverpool John Moores University, UNITED KINGDOM

Received: July 28, 2020

Accepted: April 22, 2021

Published: May 7, 2021

Copyright: © 2021 Rosnati, Fortuin. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: The data set is MIMIC-III which is freely available at <https://mimic.physionet.org/>. Our new sepsis labels for the data are available at <https://github.com/mmr12/MIMIC-III-sepsis-3-labels>.

Funding: This project was supported by the grant #2017-110 of the Strategic Focus Area 'Personalized Health and Related Technologies' of the ETH Domain (www.sfa-phrt.ch) for the SPHN/PHRT Driver Project 'Personalized Swiss Sepsis Study'. MR was supported by the Grant No. P/S023283/1 of the UKRI CDT in AI for Healthcare

Abstract

With a mortality rate of 5.4 million lives worldwide every year and a healthcare cost of more than 16 billion dollars in the USA alone, sepsis is one of the leading causes of hospital mortality and an increasing concern in the ageing western world. Recently, medical and technological advances have helped re-define the illness criteria of this disease, which is otherwise poorly understood by the medical society. Together with the rise of widely accessible Electronic Health Records, the advances in data mining and complex nonlinear algorithms are a promising avenue for the early detection of sepsis. This work contributes to the research effort in the field of automated sepsis detection with an open-access labelling of the medical MIMIC-III data set. Moreover, we propose MGP-AttTCN: a joint multitask Gaussian Process and attention-based deep learning model to early predict the occurrence of sepsis in an interpretable manner. We show that our model outperforms the current state-of-the-art and present evidence that different labelling heuristics lead to discrepancies in task difficulty. For instance, when predicting sepsis five hours prior to onset on our new realistic labels, our proposed model achieves an area under the ROC curve of 0.660 and an area under the PR curve of 0.483, whereas the (less interpretable) previous state-of-the-art model (MGP-TCN) achieves 0.635 AUROC and 0.460 AUPR and the popular commercial InSight model achieves 0.490 AUROC and 0.359 AUPR.

Introduction

Every year, it is estimated that 31.5 million people worldwide contract sepsis. With a mortality rate of 17% in its benign state and 26% for its severe state [1], sepsis is one of the leading causes of hospital mortality [2], costing the healthcare system more than 16 billion dollars in the USA alone [3]. Studies demonstrated that early treatment has a significant positive effect on the survival rate [4, 5]. In particular, [6] demonstrated that each hour delay in treating a patient results in a 7.6% increase in mortality.

Current methods of screening, such as the Modified Early Warning System (MEWS) and the Systemic Inflammatory Response Syndrome (SIRS) have been criticised for their lack of specificity, leading to low accuracies and high false alarm rates. In 2015, the Third International Consensus Definitions for Sepsis [7–9] committee worked towards incorporating

(<http://ai4health.io>). VF was supported by a PhD fellowship from the Swiss Data Science Center (www.datascience.ch). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

medical and technological advances into an up-to-date definition of sepsis, providing scientists with widely acknowledged illness criteria. Together with the rise of Electronic Health Records (EHR), the scientific community is now armed with both the data and labelling techniques to experiment with novel prediction methods [10–14], which are already proving effective in increasing survival rate [15] and promising in decreasing costs.

The models developed so far either relied on some interpretable yet simple prediction methods, such as logistic regression [13] and decision tree classifiers [16, 17], or on effective yet black-box methods such as Recurrent Neural Networks [18]. Moreover, the results achieved by different authors are rarely comparable: although most use the MIMIC-III data set, the disparities in labelling rules result in highly variable data sets (eg. [19] have 17,898 septic patients vs. 2,577 for [14]).

This work presents an attempt at reconciling interpretability and predictive performance on the sepsis prediction task and makes the following contributions:

- Gold standard for labelling. We provide a gold standard for Sepsis-3 labelling implemented on the MIMIC-III data set.
- Novel interpretable model. We present an explainable and end-to-end trainable model based on Multitask Gaussian Processes and Attentive Neural Networks for the early prediction of sepsis.
- Empirical evaluation. We assess our model on real-world medical data and report superior predictive performance and interpretability compared to previous methods.

An overview of our proposed method is shown in Fig 1. The code for labeling the data <https://github.com/mmr12/MIMIC-III-sepsis-3-labels> and for running the models <https://github.com/mmr12/MGP-AttTCN> is publicly available.

Related work

Medical time series diagnostics

Multiple researchers have tackled the task of predicting sepsis and septic shock. Works on septic shock include exploration of survival models [11] and Hidden Markov Models [12]. However, these models rely on the assumption that a patient has already developed sepsis and focus on patterns of patients' further deterioration. Other authors [13, 14, 16, 17] use linear models and decision trees on engineered features to predict sepsis onset, thus failing to capture temporal patterns. More recently, [19, 20] employed recurrent neural networks to better capture time dependencies. Crucially, all these models rely on either averaging or forward imputation of data points to create equidistant inputs. This transformation creates data artefacts and discards relevant uncertainty: in the medical field, the absence of data is a conscious decision

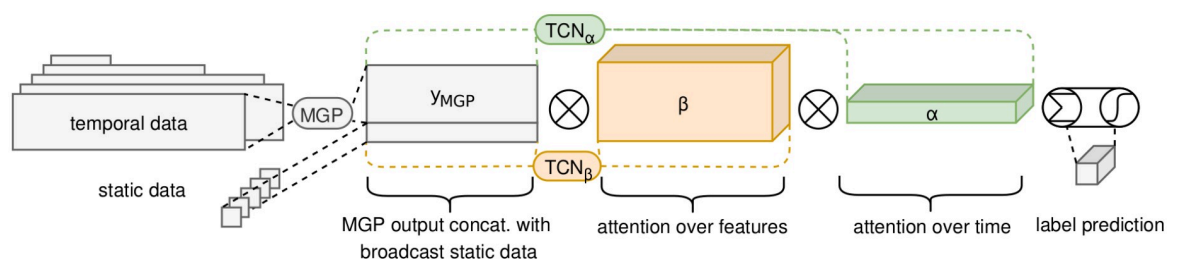


Fig 1. Proposed model architecture.

<https://doi.org/10.1371/journal.pone.0251248.g001>

made by professionals implying an underlying belief of the patient state. [18, 21] tackled this issue with Multitask Gaussian Processes (MGPs), however their models lack the interpretability necessary in the medical field.

Irregularly sampled time series

The most common solution to missing values is forward imputation [13]. [22] utilise forward imputation coupled with a missingness indicator fed into a black-box model. Although this method retains information about data presence, it is not clear how the information is later interpreted by the model and hence does not meet our transparency criteria. [23] use MGPs to fit sparse medical data, however they optimise their model for the data fit and use the parameterisation as input for a classifier rather than optimising the model for a classification task. Both [21, 24] use MGPs with end-to-end training, although their temporal covariance function is shared across all variables. Finally, [18] uses MGPs with multiple time kernels in a similar fashion to our model, although they infer the number of kernels from hyperparameter tuning rather than the data itself.

Attention based neural networks

Attention was first introduced on the topic of machine translation [25]. Since then, the concept has been used in natural language processing [26, 27] and image analysis [28, 29]. In the same spirit, [30] used attention mechanisms to improve the performance of a time series prediction model. Although their model easily explains the variable importance, its attention mechanism is based on Long Short Term Memory encodings of the time series. At any given time, such an encoding contains both the information of the current time point and all previous time points seen by the recurrent model. As such, the time domain attention does not allow for easy interpretation. More similar to our implementation is the RETAIN model [31], which generates its attention weights through reversed recurrent networks and applies them to a simple embedding of the time series. The model employs recurrent neural networks which are slower to train and suffer from the vanishing gradient problem. Furthermore, the initial and final embeddings decrease the model's interpretability. Attention in combination with a Temporal Convolutional Network (TCN) has also been used by [32], but there it only attends to time points and not to different features.

Parallel work on sepsis prediction

As mentioned above, sepsis prediction on the ICU is an important and timely problem and an active area of research. Under these circumstances, it is not surprising that some approaches have been developed in parallel to this work [33–43]. It will be an exciting and important avenue for future work to benchmark all these approaches (including ours) against each other and to compare their performances on a unified and realistic set of sepsis labels, for instance, the ones we propose in this work.

Method

In the following, we will provide a detailed explanation of our proposed model and its different components. A graphical overview of the model is shown in Fig 1.

Notation

Let us first define some notation for the problem at hand. For each patient encounter p , several features $y_{p,ti}$, k are recorded at times $t_{p,k,i}$ from admission, where $k \in \{1, \dots, M\}$ is the feature

identifier. These features are often vital signs and laboratory results. As such, they are rarely observed at the same times. Hence, we have a sparse matrix representation of observations

$$\begin{pmatrix} y_{p,1,t_1} & \cdots & y_{p,1,t_{N_p}} \\ \vdots & \ddots & \vdots \\ y_{p,M,t_1} & \cdots & y_{p,M,t_{N_p}} \end{pmatrix} \tag{1}$$

where N_p is the patient’s observation period length. We also define static features $\mathbf{s}_p = \{s_{p,M+1}, \dots, s_{p,M+Q}\}$ with feature identifiers $k \in \{M + 1, \dots, M + Q\}$, corresponding to time-independent quantities, such as age, gender and first admission unit. Finally, we define sepsis labels $l_p \in \{0, 1\}$. Given the sparsity of the data, we can define the compact representation of all observed values:

$$\{\mathbf{t}_p, \mathbf{y}_p, \mathbf{s}_p, l_p\} = \left\{ \begin{array}{l} \{t_{p,i,k}\}_{i \in \{0, \dots, N_p\}, k \in \{1, \dots, M\}}, \\ \{y_{p,i,k}\}_{i \in \{0, \dots, N_p\}, k \in \{1, \dots, M\}}, \\ \{s_{p,M+1}, \dots, s_{p,M+Q}\}, \\ l_p \end{array} \right\} \tag{2}$$

The goal of the model is, for a given set $\{\mathbf{t}_p, \mathbf{y}_p, \mathbf{s}_p\}$ to predict the label l_p . In order to remove clutter, we will from now on drop the patient-specific subscript p from all notation, and the feature subscript k from time notation, simplifying $t_{p,k,i}$ to t_i .

Multitask Gaussian Process (MGP)

Gaussian processes are non-parametric Bayesian models commonly known for their ability to generate coherent function fits to a set of irregular samples, by modelling the data covariance. As they easily account for uncertainty and do not require homogeneously sampled data, Gaussian processes are the perfect candidate model to deal with irregularly sampled medical time series.

We use a Multitask Gaussian Process (MGP) [44] to capture feature correlation and [45]’s end-to-end training framework, in a similar manner to [24]. Given an hourly spaced time series $\{t'_i\}_{i=-N_p}^0$ where 0 is the time of prediction, the MGP layer produces a set of posterior predictions for each feature, which will then be fed into a classification model.

We define a patient-independent prior over the true values of $\{y_{i,k}\}$ by $\{f_k(t_i)\}$ such that

$$y_{i,k} \sim \mathcal{N}(f_k(t_i), \sigma_k^2) \tag{3}$$

$$\langle f_k(t_i), f_{k'}(t_j) \rangle = \sum_{l \in L} K_l^k(k, k') K_l^{tt}(t_i, t_j) \tag{4}$$

where $\{K_l^{tt}(t_i, t_j)\}_{l \in L}$ are parametric time point covariances varying in smoothness, $\{K_l^k(k, k')\}_{l \in L}$ are free-form feature covariances at a given smoothness level, independent of time, and L are smoothness clusters. Over all variables and time points, the multivariate model has covariance

$$\sum_{l \in L} K_l^k \otimes K_l^{tt} + D \otimes I \tag{5}$$

where $D = \text{diag}(\sigma_k)$ are the noise terms associated to each feature and \otimes is the Kronecker

product. This formulation allows each datapoint to be defined as both a function of its own timed observations and observations of the remaining features.

The quick illness progression is well suited to the quadratic growth of the covariance matrix. However, for the few cases when a patient develops sepsis well into their hospitalisation, a suitable measure to prevent excessive memory consumption is to ignore the initial datapoints of the patient.

Note that there are two main feature clusters: vital signs (vitals) and laboratory results (labs). Vitals are noisier and sampled more often, whereas labs are more monotone and rarely sampled. As opposed to [18], we do not treat the number of clusters L as hyperparameters but set $L = 2$ and define

$$K_i^t(t_i, t_j) = \exp\left(\frac{-|t_i - t_j|}{\lambda_i}\right) \tag{6}$$

as Ornstein-Uhlenbeck (OU) kernels with lengths λ_1 and λ_2 , each representing a cluster smoothness. OU kernels are well suited to capture local variations and do not assume infinite differentiability as Squared Exponential kernels do. In our case, differentiability implies a level of smoothness which is unrealistic for medical records and only introduces unnecessary bias. In addition, given the scarce availability of labs, short lengthscales would be an ill fit to the data. We hence discarded kernels varying over lengthscales such as the Cauchy and the Rational Quadratic kernels. $K_i^k(k, k')$ are free-form covariance matrices that are learned by gradient descent.

The posterior over the regularly sampled timepoints $\mathbf{t}' = \{t'_i\}_{i=-N_p}^0$ is a multivariate Gaussian with mean and covariance:

$$\begin{aligned} \mu &= \left(\sum_{l \in L} K_l^k \otimes K_l^{t' t'}\right) \left(\sum_{l \in L} K_l^k \otimes K_l^{t t} + D \otimes I\right)^{-1} \mathbf{y} \\ \Sigma &= \sum_{l \in L} K_l^k \otimes K_l^{t' t'} \\ &\quad - \left(\sum_{l \in L} K_l^k \otimes K_l^{t' t'}\right) \left(\sum_{l \in L} K_l^k \otimes K_l^{t t} + D \otimes I\right)^{-1} \left(\sum_{l \in L} K_l^k \otimes K_l^{t' t}\right) \end{aligned} \tag{7}$$

In order to approximate the posterior distribution, we then take Monte Carlo samples \mathbf{y}_{MC} from $\mathbf{Y}_{MGP} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$.

To feed the MGP samples into the classifier, we fix the model time window to $N = 25$ by either zero padding or truncating the beginning of the time series. We choose to do so at the beginning of the time series in order to align prediction times as the last step of the temporal classification model. Here, we also integrate the static variables by broadcasting them over each time. The reasoning behind this design choice is explained in more details in the following section.

Attention Time Convolutional Network (AttTCN)

The concept of attention was born in machine translation [25] and has recently successfully been applied to different types of sequential data [26–30]. In machine translation, given an input sentence embedding

$$S = \{\mathbf{h}_1, \dots, \mathbf{h}_{|S|}\} \tag{8}$$

the attention mechanism produces weights

$$\{\alpha_1^i, \dots, \alpha_{|S|}^i\} \quad \text{s.t.} \quad \alpha_j^i \in [0, 1], \quad \sum_j \alpha_j^i = 1 \tag{9}$$

and a context vector

$$\mathbf{c}_i = \sum_j \alpha_j^i \mathbf{h}_j \tag{10}$$

used to predict target word i . The weights α_j^i can therefore be interpreted as the importance of the input sentence's j^{th} word to produce the i^{th} word of the translation.

More recently, [31] have applied attention to clinical time series. Given a time series

$$\{\mathbf{x}_1, \dots, \mathbf{x}_T\} \subset \mathbb{R}^r, \tag{11}$$

the authors first create a time-independent embedding of the data

$$\{\mathbf{v}_1, \dots, \mathbf{v}_T\} \subset \mathbb{R}^m. \tag{12}$$

They then use inversed recurrent neural networks (RNN) to create weights $\boldsymbol{\alpha} \in \mathbb{R}^T$ and $\boldsymbol{\beta} \in \mathbb{R}^{T \times m}$, where $\alpha_j \in [0, 1]$ and $\beta_{ij} \in [-1, 1]$, with softmax and tanh activations respectively. The context vectors then take the form

$$\mathbf{c}_i = \sum_{j \leq i} \alpha_j \boldsymbol{\beta}_j \odot \mathbf{v}_j, \tag{13}$$

where \odot is the element-wise product, and are fed into a multilayer perceptron with softmax activation to yield a prediction.

The attention model we devised borrows some ideas from [31]. Two embeddings, $\mathbf{z} = [z_1, \dots, z_N]$ and $\mathbf{z}' = [z'_1, \dots, z'_N]$ with $\mathbf{z}, \mathbf{z}' \in \mathbb{R}^{N \times (M+Q)}$, are directly generated from the interpolated data $\mathbf{y}_{MC} \in \mathbb{R}^{N \times (M+Q)}$ through two temporal convolutional networks (TCNs), removing a layer of abstraction and hence facilitating interpretability.

TCNs are a class of neural networks composed of *causal* convolutions stacked into Residual Blocks. A causal convolution is a 1D convolutional layer which only takes inputs from the past to generate its output, avoiding any information leakage from the future. Residual Blocks are made of two causal convolutional layers together with ReLU activation functions, dropout and L2 regularisations. The Residual Blocks also include an identity map from the input of the block added to the output. As we only use up to 12 layers, this last step is omitted in our architecture. TCNs have shown to outperform RNNs [46], are faster at training and do not suffer from vanishing gradients. Given the latter, inverting the time series similarly to [47] also becomes an unnecessary step which we omit.

We generate the attention weights α and β as

$$\alpha_{j,0} = \text{softmax}(\mathbf{z}_j \times \mathbf{W}_{\alpha,0} + \mathbf{b}_{\alpha,0}) \tag{14}$$

$$\alpha_{j,1} = \text{softmax}(\mathbf{z}_j \times \mathbf{W}_{\alpha,1} + \mathbf{b}_{\alpha,1}) \tag{15}$$

$$\boldsymbol{\beta}_{j,0} = \text{sigmoid}(\mathbf{z}'_j \times \mathbf{W}_{\beta,0} + \mathbf{b}_{\beta,0}) \tag{16}$$

$$\boldsymbol{\beta}_{j,1} = \text{sigmoid}(\mathbf{z}'_j \times \mathbf{W}_{\beta,1} + \mathbf{b}_{\beta,1}) \tag{17}$$

$$\mathbf{W}_{\alpha,0}, \mathbf{W}_{\alpha,1} \in \mathbb{R}^{M+Q} \quad b_{\alpha,0}, b_{\alpha,1} \in \mathbb{R} \tag{18}$$

$$\mathbf{W}_{\beta,0}, \mathbf{W}_{\beta,1} \in \mathbb{R}^{(M+Q) \times (M+Q)} \quad \mathbf{b}_{\beta,0}, \mathbf{b}_{\beta,1} \in \mathbb{R}^{M+Q}, \tag{19}$$

such that $\boldsymbol{\alpha} = [\boldsymbol{\alpha}_0, \boldsymbol{\alpha}_1] \in \mathbb{R}^{N \times 2}$ and $\boldsymbol{\beta} = [\boldsymbol{\beta}_0, \boldsymbol{\beta}_1] \in \mathbb{R}^{N \times (M+Q) \times 2}$.

We then create two context vectors, one for each of the negative and positive label predictions

$$\mathbf{c}_i = \sum_{j \leq i} \alpha_{j,\delta} \boldsymbol{\beta}_{j,\delta} \odot \mathbf{y}_{MC,j} \in \mathbb{R}^{N \times (M+Q) \times 2}, \quad \delta \in \{0, 1\}, \tag{20}$$

where $\mathbf{y}_{MC,j}$ is broadcast to meet the dimensionality of $\boldsymbol{\beta}_{j,\delta}$. We then predict the labels as

$$\hat{\mathbf{I}}_i = \text{softmax}\left(\sum_n \sum_m \mathbf{c}_{i,nm}\right) \in [0, 1]^2. \tag{21}$$

In our case, we are only interested in making predictions with the latest available data. We therefore only use $\hat{\mathbf{I}}_{\text{last}}$ to train the model. This of course can be easily modified to suit any specific use case.

Since the MGP output is directly multiplied by weights c_i , the classification model can be interpreted as a scoring mechanism where each past point $y_{MC,ij}$ contributes $\alpha_{i,0} \beta_{ij,0}$ to the time series being classified as positive, and $\alpha_{i,1} \beta_{ij,1}$ to the time series being classified as negative. The positive and negative scores are then normalised to represent probabilities of the positive or negative labelling. As we designed both $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ to be non-negative, we can hence directly look at the average $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ over Monte Carlo samples to see which time points and features contribute most strongly to the network’s positive or negative decision. This facilitates the interpretability of our model compared to previous approaches.

Data

Sepsis is defined as a life-threatening organ dysfunction caused by a dysregulated host response to infection [7]. A dysregulated host response is interpreted as a suspicion of infection. In EHR terms, it is encoded by the administration of high spectrum antibiotics and a bacterial blood culture within a set interval of each other. The organ dysfunction is interpreted as a two point increase in Sequential Organ Failure Assessment (SOFA) within a suspected infection window. The SOFA score quantifies the deterioration of different systems—respiration, coagulation, liver function, cardiovascular function, central nervous system, and renal function.

We make use of the MIMIC-III dataset, a collection of medical records for over 40’000 patients who stayed in critical care units of the Beth Israel Deaconess Medical Center between 2001 and 2012 [48]. The records are composed of vital sign recordings, laboratory tests, drugs administered, and patients’ outputs. We encode the Sepsis-3 criteria in the MIMIC-III dataset following [48, 49]’s code available on GitHub and [21]’s code that the authors have generously provided.

One key difference between our assumptions and the ones [21] develop is the handling of missing SOFA contributor values: if one or more SOFA contributors are missing, [21] do not calculate the total score. On the other hand, we assume such a contributor to be within a healthy norm, hence implying a zero contribution. With our methodology, patients worsening in one area but with no measurements in another will be considered septic, whereas they will

not in the [21] dataset. After discussions with a clinician, we learned that the standard practice in the clinic would be to assume healthy values for all unmeasured variables, as we did in our labeling. Moreover, if the treating physician would expect a variable to be outside the healthy norm, they would usually measure it, such that most unmeasured variables will indeed have a high probability of being healthy. Our labeling approach thus fits better with the standard clinical practice and includes more septic patients that would not be included in [21]’s data. Moreover, we hypothesize that these patients would be the ones where the treating clinicians do not already suspect a sepsis and have thus not measured all the SOFA variables, which will potentially make them harder to classify but arguably also more interesting, because they would otherwise be missed by the treating doctors.

In order to validate our results, we carry out all experiments on using both labelling techniques.

Patient inclusion

We filter for patients admitted to Intensive Care Units (ICU) who are more than 14 years old and with valid records. Case patients are patients having sepsis onset within their ICU stay, whereas control patients have not developed sepsis nor have an ICD discharge code referring to sepsis. Starting with 58,976 patients, we find 14,071 control patients and 7,936 case patients using our labels, versus 1,797 cases using [21]’s labels.

Feature extraction

Reviewing sepsis-related literature and commonly extracted laboratory and vital recordings, we extract all features which were reported at least once for more than 75% of the included population. The final 24 dynamic features are reported in Table 1. We also extract static features—age, gender, and first ICU admission department.

Case-control matching

As the goal is to predict sepsis prior to onset, the cases data was extracted between ICU admission and sepsis onset. Note that sepsis onset happens early within ICU admission, with the median patient getting sick at 3.4 hours after admission. On the other hand, patients not developing sepsis are more likely to recover completely, and do so in a lengthier time frame. In addition, once they are close to discharge, their vitals and labs are within the norms. Hence, both the length and the values of the time series are strong discriminatory factors which ease the classification. We hence carry out a matching strategy similar to [21]: following the class imbalance ratio, we associate each control time series to a case time series and truncate the

Table 1. List of dynamic features.

Vitals	Labs	
Systolic blood pressure (sysbp)	bicarbonate	Pulse transit time (ptt)
Diastolic blood pressure (diabp)	creatinine	International normalized ratio (inr)
Mean blood pressure (meanbp)	chloride	Prothrombin time (pt)
Respiratory rate (resprate)	glucose	sodium
Heart rate (heartrate)	hematocrit	Blood urea nitrogen (bun)
Pulse oximetry (spo2 pulsoxy)	hemoglobin	White blood cells (wbc)
Body temperature (tempc)	lactate	magnesium
	platelet	Blood gas pH (ph bloodgas)
	potassium	

<https://doi.org/10.1371/journal.pone.0251248.t001>




Fig 2. Baseline patient data for different horizons. x-axis: Time from admission. y-axis: Feature identifier.

<https://doi.org/10.1371/journal.pone.0251248.g002>

control to have the same length as the case from ICU admission. We then discard patients with less than 40 data points within the selected window, and—for computational tractability—truncate the first $N_p - 250$ initial values of patients' time series in order to keep a maximum of 250 data points per patient.

Horizon augmentation

As our goal is to predict sepsis early, we augment the data by creating new shorter time series. For each time series, we create six copies, where each copy represents a different horizon to onset. We then proceed to truncate the last one to six hours prior to onset from the time series copies. In order to keep data consistency, we once again discard time series with less than 40 observations. Fig 2 is a graphical representation of the discretised version used for the baseline of an augmented datapoint, whereas in Tables 2 and 3 we illustrate the data distribution per horizon.

Table 2. Augmented dataset statistics with Moor et al. [21] labels.

Horizon to onset	N. of patients	N. of obs. per patient
0 h	15,123	69.9±59.6
1 h	11,258	56.6±59.1
2 h	8,478	61.4±62.8
3 h	6,554	66.5±65.9
4 h	5,233	70.6±69.0
5 h	4,162	76.3±71.9
6 h	3,390	81.9±74.3

<https://doi.org/10.1371/journal.pone.0251248.t002>

Table 3. Augmented dataset statistics with our labels.

Horizon to onset	N. of patients	N. of obs. per patient
0 h	20,075	64.0 ± 65.5
1 h	15,832	62.5 ± 65.6
2 h	12,080	66.1 ± 67.2
3 h	9,441	69.7 ± 68.1
4 h	7,484	73.4 ± 68.5
5 h	6,007	77.1 ± 68.2
6 h	4,876	81.2 ± 67.4

<https://doi.org/10.1371/journal.pone.0251248.t003>

Data split

Finally, we split the data into training, validation, and testing sets, respectively capturing 80%, 10%, and 10% of the data. We then normalise the data by subtracting the training set mean and dividing by the training set standard deviation of each feature.

Hyperparameter search and training

As the datasets are highly imbalanced, we carry out a case set oversampling: we randomly resample the case set to have the same size as the control set. In addition, at each iteration we sample the same number of cases and controls and then feed a shuffled version into the model. In this manner, the model will see an equal number of controls and cases and will not become biased towards zero labels. This procedure is not applied to the validation and test sets, as the results would not compare to real-life settings.

For both our model MGP-AttTCN and all baselines, in order to select the best possible hyperparameters, we performed a hyperparameter random search, as described in Table 4.

Baselines

Data preparation

In order to benchmark our MGP model, we build some baselines homogenising the data sampling. For each hour and variable, we take the average of the available observations. If a given hour has no observations, we carry forward the average of the previous hour. In this manner, we generate an hourly sampled time series for each patient. We then proceed to normalise the size of each patient matrix by setting a time window of observation N . For patients having more than N observations N_p , we discard the first $N - N_p$ observation; whereas for patients having less than N observations, we pad the beginning of the matrix with zeros.

$$\begin{pmatrix} y_{p,1,t_1} & \cdots & y_{p,1,t_{N_p}} \\ \vdots & \ddots & \vdots \\ y_{p,M,t_1} & \cdots & y_{p,M,t_{N_p}} \end{pmatrix} \xrightarrow{\text{carry forward}} \begin{pmatrix} y_{p,1,1} & \cdots & y_{p,1,N_p} \\ \vdots & \ddots & \vdots \\ y_{p,M,1} & \cdots & y_{p,M,N_p} \end{pmatrix} \tag{22}$$

Table 4. Hyperparameter search.

Hyperparameter	Random Search	
	min	max
MGP Monte Carlo samples	4	20
TCN kernel size	2	6
TCN number of Residual Blocks	2	12
TCN number of hidden layers	10	55
TCN dropout rate	0	0.99
TCN L2 regularisation	0	250

<https://doi.org/10.1371/journal.pone.0251248.t004>

$$\xrightarrow{\text{normalise}} \left\{ \begin{array}{l} \begin{pmatrix} y_{p,1,N-N_p} & \cdots & y_{p,1,N} \\ \vdots & \ddots & \vdots \\ y_{p,M,N-N_p} & \cdots & y_{p,M,N} \end{pmatrix} \text{ if } N_p \geq N \\ \begin{pmatrix} 0 & \cdots & 0 & y_{p,1,N-N_p} & \cdots & y_{p,1,N} \\ \vdots & & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & y_{p,M,N-N_p} & \cdots & y_{p,M,N} \end{pmatrix} \text{ else} \end{array} \right. \quad (23)$$

We choose to align the end of the time series as opposed to the beginning, as the relative importance of time points is higher closer to when a patient becomes sick rather to when they are admitted to the ICU.

As a next step, we augment the data to focus on different time series in a similar manner as for irregularly sampled data. We create seven copies of each time series, where for each copy we discard the last zero to six hours, then normalise the matrix as above. We hence generate a dataset $\mathbf{Y}_{BL} = \{Y\}_q = \{\{y_{q,ij}\}_{i,j=1}^{N,M}\}_q$ where q represents all augmented the time series.

InSight

The *InSight* scoring model is one of the few machine learning algorithms to surpass the proof-of-concept stage with multiple research, economic and clinical trials [13, 14, 16, 50]. We therefore include it as a baseline to our model. The key concept of the model is to use few largely available vitals, build some handcrafted features and train a simple classification model.

Here we provide a brief summary of the main method. The features extracted are based on a window of six consecutive hours. For each six hour window, we extract each variable’s mean M_i and difference D_i (last observation minus first observation) over the window. We also extract variable pairs correlation D_{ij} and *triplet correlation* D_{ijk} , where i, j, k are observed variables. We interpret the latter as a relaxation of the Pearson correlation: if the correlation between two variables is

$$\rho_{XY} = \frac{\mathbb{E}[(X - \mu_X)(Y - \mu_Y)]}{\sigma_X \sigma_Y} \quad (24)$$

then we define the *triplet correlation* as

$$\rho_{XYZ} = \frac{\mathbb{E}[(X - \mu_X)(Y - \mu_Y)(Z - \mu_Z)]}{\sigma_X \sigma_Y \sigma_Z} \quad (25)$$

We then classify the difference and correlations as either positive, negligible or negative using their distribution quantiles over every patient and six hour window observed. Note that given the high level of data missingness, many variables are calculated by forward imputation and hence have no variance over six hours. To adjust for the high number of zero correlations, we

calculate the quantiles of non-zero correlations and define:

$$\hat{D}_i = \begin{cases} 1 & \text{if } D_i > q^*(2/3) \\ -1 & \text{if } D_i < q^*(1/3) \\ 0 & \text{otherwise} \end{cases} \quad (26)$$

where q^* is the adjusted quantile function. We proceed in a similar manner for the correlations and triplet correlations.

In order to keep the results comparable to the AttTCN fixed window N , we extract $N - (6 - 1)$ six consecutive hour window and vectorise the resulting features, generating in total

$$n_{\text{features}} = (N - 5) \times (2 \times M + \binom{M}{2} + \binom{M}{3}) \quad (27)$$

features per patient.

To remain consistent with the original work, we only kept patients with at least one observation for each feature over the 5 hour period for the following observations: age, systolic blood pressure, pulse transit time, heart rate, temperature, respiration rate, white blood cell count, pH and pulse oximetry. The corresponding dataset statistics can be found in Tables 5 and 6.

Although the original paper does not specify which classification method the authors employ, we derive by their description of a *dimensionless score* that the method is a logistic regression.

Other baselines

Logistic regression. As a simple baseline, we perform a Ridge Logistic Regression using the hourly data described above.

Table 5. Baseline dataset with our labels.

Horizon to onset	N. of patients	N. of patients with onset
0 h	1,943	786 (40.5%)
1 h	1,938	782 (40.4%)
2 h	1,810	690 (38.1%)
3 h	1,657	606 (36.6%)
4 h	1,504	532 (35.4%)
5 h	1,352	467 (34.5%)
6 h	1,217	420 (34.5%)

<https://doi.org/10.1371/journal.pone.0251248.t005>

Table 6. Baseline dataset with Moor et al. [21] labels.

Horizon to onset	N. of patients	N. of patients with onset
0 h	2,298	658 (28.6%)
1 h	2,294	657 (28.6%)
2 h	2,187	609 (27.8%)
3 h	1,981	531 (26.8%)
4 h	1,748	453 (25.9%)
5 h	1,518	384 (25.3%)
6 h	1,330	329 (24.7%)

<https://doi.org/10.1371/journal.pone.0251248.t006>

Ablation models. In addition to Insight and the Logistic Regression, we perform ablation studies on our proposed model. In a first instance, we remove the AttTCN arm and replace it by a logistic regression (model “MGP-log.reg.”). Secondly, we remove the arm of the AttTCN controlling the attention over time α (model “MGP-AttTCN w/o α ”), then the arm generating β (model “MGP-AttTCN w/o β ”).

MGP-TCN. Finally, we compare our model to the state-of-the-art MGP-TCN algorithm [21].

Experimental results

We compare our model’s performance to the performance of the InSight algorithm [13] and to the state-of-the-art MGP-TCN algorithm [21]. Fig 3 shows the predictive performance of the models for different time horizons, whereas the numerical results can be found in the Appendix, in S1–S8 Tables in [S1 File](#).

Comparison between different data labels

The first result is the difference in performance of models applied to the different labelling methods. The SOFA contributor assumption from [21] has two main implications. Firstly, it considerably restricts the number of patients. Assuming that sicker patients receive more medical attention, the patients included are likely to be in worse conditions than the septic patients excluded and hence easier to classify. Secondly, it delays sepsis onset. For example, a patient having a severe liver failure with few other recorded vitals, followed by an overall collapse further in time will have septic onset at the time of its liver failure in our records, whereas it will only be considered septic at the time of the overall collapse in [21]’s labels. On the other hand, the labels we produce reflect the incomplete nature of medical data: even if only a part of all the potentially relevant tests are carried out at any given time, a doctor must be able to assess a patient’s well-being and foresee potential complications. The difference in labels implies a discrepancy in task difficulty: [21]’s labels present an easier learning problem, but define a more narrow use case in real-world scenarios.

Indeed, when assessing the performance of the different models on the two different data labellings, it becomes evident that our proposed labels are harder to fit. This means that predicting sepsis in a realistic setting on the intensive care unit is probably much harder than previous work would suggest.

Model performance

We find that our MGP-AttTCN model yields a better performance than the MGP-TCN [21] when presented with patients further in time from sepsis onset (i.e., earlier in time) (Fig 3, top row). In the case of our labels the difference is clearly noticeable, whereas with [21]’s labels it is of lower statistical significance.

Moreover, we observe that ablations of our model (e.g., changing the GP kernel, removing the α weights, or removing the β weights) reduces our model’s performance, as expected (Fig 3, middle row). The exception to this are the α weights on our labels, which seem to neither improve nor impair our model’s performance significantly. Note however, that the α and β weights play a strong role for the interpretability of our model (see below) and are thus useful even without influencing the raw predictive performance.

Finally, with our labels, our model also outperforms InSight, as well as the simple logistic regression and MGP-LogReg baselines (Fig 3, bottom row). The intuition behind this result is the robustness of the model to missing data: It accounts for the data uncertainty and hence has a better performance on lower resolution and more irregular data. Note however that on [21]’s

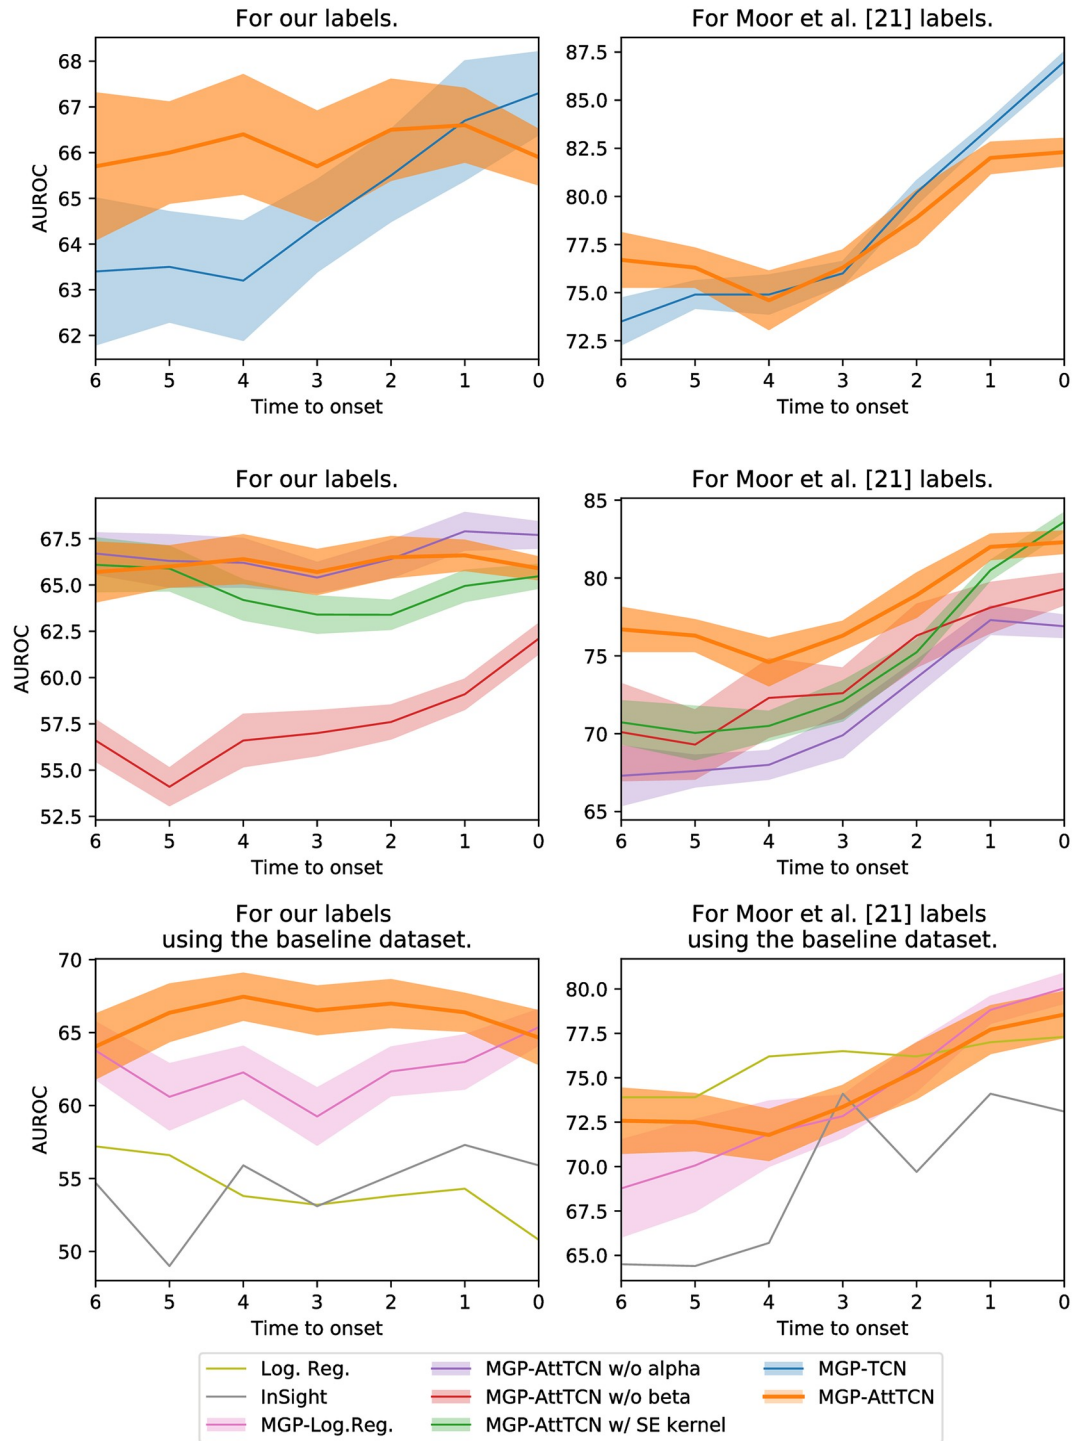


Fig 3. Area under the ROC curve of different models. It can be seen that our proposed labels are harder to fit than the ones by [21]. Moreover, our proposed model outperforms the baselines on both label sets, especially for earlier prediction horizons.

<https://doi.org/10.1371/journal.pone.0251248.g003>




Fig 4. Heatmaps of the learned MGP covariance matrices between the data features for the two different smoothness clusters.

<https://doi.org/10.1371/journal.pone.0251248.g004>

labels, the logistic regression is a stronger competitor (which was not considered in their paper), highlighting again that their prediction task is significantly easier than the one with our more realistic labels. All these results were measuring the performance using the area under the receiver-operator characteristic curve (AUROC), but we provide additional results using the precision-recall curve (AUPRC) in the appendix (S1 Fig in [S1 File](#)), which qualitatively show the same observations.

MGP interpretability

Inspecting the learned covariances ([Fig 4](#)), we notice that the two OU lengthscales converged to represent two clusters within the selected variables: a shorter lengthscales (around two hours) represents noisy data, whereas a larger lengthscales (around 64 hours) represents smoother observations. In addition, the feature covariance matrix for the short lengthscales puts more emphasis on vitals, while the one for the long lengthscales puts more emphasis on labs, fitting our initial intuition that vitals vary more rapidly. Graphically, one can observe this by inspecting the diagonals on the covariance heatmaps.

On a more granular level, the two covariance matrices also provide insights about the underlying variables. One can for instance observe that the body temperature (*tempc*) has a larger variance than the systolic and diastolic blood pressure (*sysbp*, *diabb*), following the general clinical intuition. Moreover, we can observe correlations between different features, such as a negative correlation between temperature and heart rate, which also seems to coincide with the general medical expectation. These covariances can then for instance be used by the model to extrapolate a full time series from a single INR observation with an inverse correlation to the pulse oximetry observations ([Fig 5](#)).

Attention weights

Even if our model does not outperform the MGP-TCN under all conditions, its main advantage over the baselines lies in its improved interpretability due to the attention mechanism. Once the samples have been drawn, the weights α and β provide us with information about

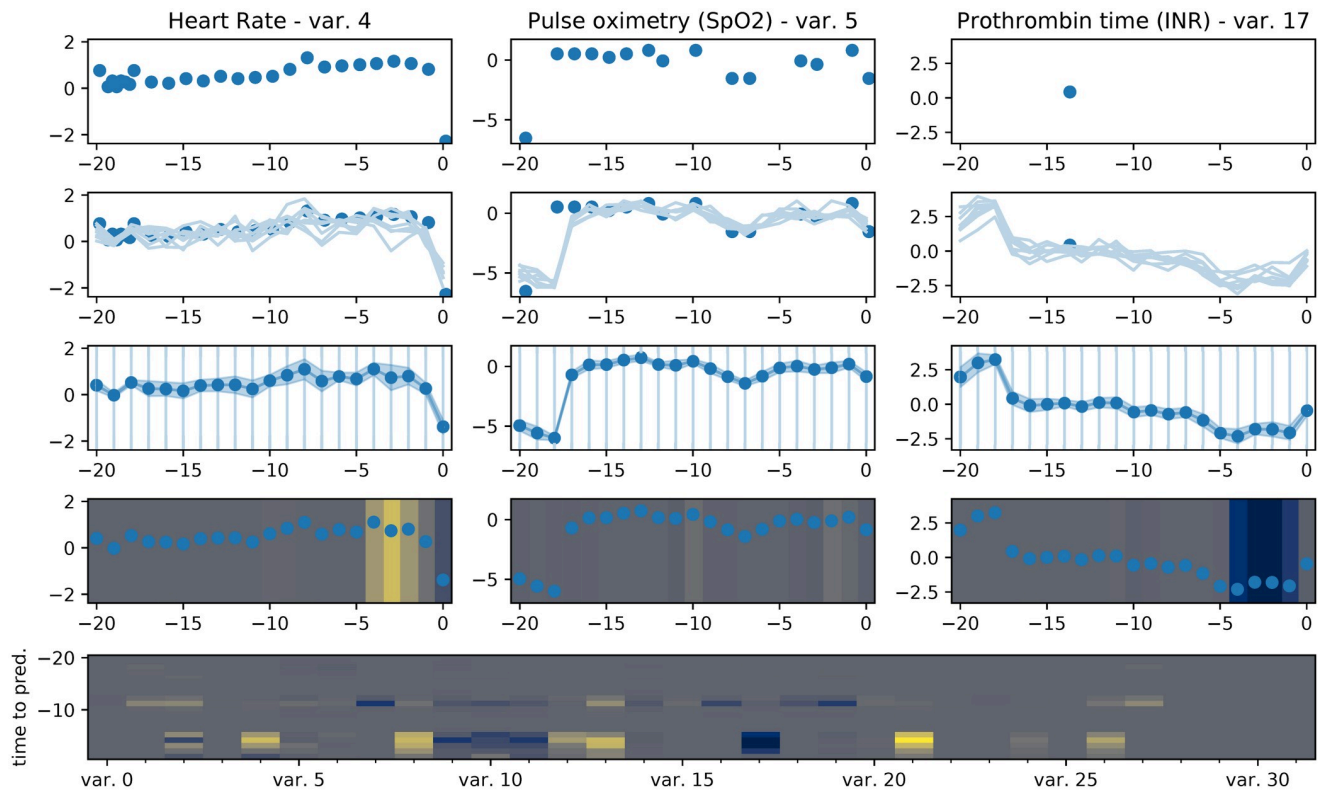


Fig 5. Visualization of the journey of an exemplary patient trajectory through our proposed model architecture. The raw features (row 1), measured at irregular time points, are interpolated by the MGP (row 2). Samples from the MGP posterior can then be aggregated into means and variances for each feature on a fixed, regularly-spaced time grid (row 3). These values are then attended to by the TCN (row 4), where positive attention weights are yellow and negative ones blue. Row 5 shows the attention weights separated by features (x-axis) and time points (y-axis).

<https://doi.org/10.1371/journal.pone.0251248.g005>

the importance of different time points and features for the model's prediction. The attention weights for an exemplary patient trajectory are depicted in Fig 5.

The figure shows the flow of data from a randomly chosen example patient through our model. In the first row, we see the actual measured data. We can see that while for instance the heart rate and oximetry are measured regularly, the prothrombin time has a lot of missing values. In the second row, we see imputations of the time series sampled from our MGP. We see that even though the prothrombin time measurements are sparse, the MGP yields imputations with reasonably low uncertainty, thanks to information extracted by the model from the other features. In the third row, we resample the time series from the MGP interpolations on a regular grid, which includes a mean and uncertainty estimate for each value. Finally, in the fourth row, the TCN part of our model can assign attention weights to the different resampled measurements, which show their influence on the model's final prediction. Positive attention weights mean that the respective feature increases the model's probability of diagnosing the patient as septic. They are shown in yellow. Negative attention weights decrease this probability and are shown in blue. The final row of the figure shows the attention weights for all the different variables (x-axis) over time (y-axis).

Overall, the absolute values of α are small for points further from the prediction time and increasingly larger closer to it. A good example of this behaviour is the fourth row in Fig 5, where feature importance increases in time. We can also see there that different features can have opposing effects on the prediction. While the elevated heart rate close to the prediction

time increases the likelihood of a sepsis prediction (first column, yellow weights), the lowered prothrombin values reduce this likelihood (third column, blue weights). These attention weights have been deemed plausible by a clinician to whom we showed the figure, demonstrating that they can help to build trust in the model's prediction by making its decision process more relatable to trained professionals and comparable with their prior knowledge. In this particular patient, for instance, the elevated heart rate is a common symptom of sepsis [51] and thus deserves its positive attention weight, while another common symptom is an *increased* prothrombin time [52], such that the *decreased* prothrombin time in this example should rightfully be regarded as speaking against the diagnosis of sepsis, as attested by its negative attention weight. Interestingly, the low prothrombin values are not actually measured in this example, but predicted by the MGP purely based on the other measured features and the learned covariances.

Finally, $\alpha \times \beta \times y_{MC}$ gives the individual score contribution of each feature at each time point. These weights are shown in the last row of the figure. It can again be seen that the attention weights are generally larger in magnitude closer to the prediction time. Moreover, about half of the features have significant non-zero attention weights, while the others seem to not be important for the prediction in this example.

These visualizations could be used by doctors to make an informed decision about whether or not to trust the prediction of the model for each given patient, thus facilitating the interpretability and accountability that is crucial in medical applications.

Conclusion

We have shown that current data sets for the early prediction of sepsis underestimate the true difficulty of the problem and proposed a new labelling for the MIMIC-III data set that corresponds more closely to a realistic intensive care setting. Moreover, we have proposed a new machine learning model, the MGP-AttTCN, which outperforms the state-of-the-art approaches on the easier labels from the literature as well as on our proposed harder labels. Additionally, our model provides an interpretable attention mechanism that will allow clinicians to make more informed decisions about trusting its predictions on a case-by-case basis.

Potential avenues for future work include a more thorough discussion with clinicians to make our proposed labels even more representative of the real-world task. Moreover, there is potential for architectural improvements, for instance by meta-learning the MGP prior [53, 54], amortizing the latent MGP inference for performance gains [55–59], discretizing the latent space for improved interpretability [60, 61], or treating the neural network parameters in a Bayesian way to improve the uncertainty estimation [62–64].

Supporting information

S1 File.
(PDF)

Acknowledgments

We would like to thank Gunnar Rätsch, Karsten Borgwardt, Michael Moor, Drago Plecko, and Nicolas Bennett for helpful discussions.

Author Contributions

Conceptualization: Vincent Fortuin.

Data curation: Margherita Rosnati.

Formal analysis: Vincent Fortuin.

Investigation: Margherita Rosnati.

Methodology: Margherita Rosnati, Vincent Fortuin.

Project administration: Vincent Fortuin.

Resources: Vincent Fortuin.

Software: Margherita Rosnati.

Supervision: Vincent Fortuin.

Validation: Margherita Rosnati, Vincent Fortuin.

Visualization: Margherita Rosnati.

Writing – original draft: Margherita Rosnati, Vincent Fortuin.

Writing – review & editing: Margherita Rosnati, Vincent Fortuin.

References

1. Fleischmann C, Scherag A, Adhikari NK, Hartog CS, Tsaganos T, Schlattmann P, et al. Assessment of global incidence and mortality of hospital-treated sepsis. Current estimates and limitations. *American journal of respiratory and critical care medicine*. 2016; 193(3):259–272. <https://doi.org/10.1164/rccm.201504-0781OC> PMID: 26414292
2. Vincent JL, Marshall JC, Namendys-Silva SA, François B, Martin-Loeches I, Lipman J, et al. Assessment of the worldwide burden of critical illness: the intensive care over nations (ICON) audit. *The Lancet Respiratory medicine*. 2014; 2(5):380–386. [https://doi.org/10.1016/S2213-2600\(14\)70061-X](https://doi.org/10.1016/S2213-2600(14)70061-X) PMID: 24740011
3. Angus DC, Linde-Zwirble WT, Lidicker J, Clermont G, Carcillo J, Pinsky MR. Epidemiology of severe sepsis in the United States: analysis of incidence, outcome, and associated costs of care. *Critical care medicine*. 2001; 29(7):1303–1310. <https://doi.org/10.1097/00003246-200107000-00002>
4. Kumar A, Roberts D, Wood KE, Light B, Parrillo JE, Sharma S, et al. Duration of hypotension before initiation of effective antimicrobial therapy is the critical determinant of survival in human septic shock. *Critical care medicine*. 2006; 34(6):1589–1596. <https://doi.org/10.1097/01.CCM.0000217961.75225.E9> PMID: 16625125
5. Nguyen HB, Corbett SW, Steele R, Banta J, Clark RT, Hayes SR, et al. Implementation of a bundle of quality indicators for the early management of severe sepsis and septic shock is associated with decreased mortality. *Critical care medicine*. 2007; 35(4):1105–1112. <https://doi.org/10.1097/01.CCM.0000259463.33848.3D> PMID: 17334251
6. Castellanos-Ortega Á, Suberviola B, García-Astudillo LA, Holanda MS, Ortiz F, Llorca J, et al. Impact of the Surviving Sepsis Campaign protocols on hospital length of stay and mortality in septic shock patients: results of a three-year follow-up quasi-experimental study. *Critical care medicine*. 2010; 38(4):1036–1043. <https://doi.org/10.1097/CCM.0b013e3181d455b6> PMID: 20154597
7. Singer M, Deutschman CS, Seymour CW, Shankar-Hari M, Annane D, Bauer M, et al. The third international consensus definitions for sepsis and septic shock (Sepsis-3). *Jama*. 2016; 315(8):801–810. <https://doi.org/10.1001/jama.2016.0287> PMID: 26903338
8. Seymour CW, Liu VX, Iwashyna TJ, Brunkhorst FM, Rea TD, Scherag A, et al. Assessment of clinical criteria for sepsis: for the Third International Consensus Definitions for Sepsis and Septic Shock (Sepsis-3). *Jama*. 2016; 315(8):762–774. <https://doi.org/10.1001/jama.2016.0288> PMID: 26903335
9. Shankar-Hari M, Phillips GS, Levy ML, Seymour CW, Liu VX, Deutschman CS, et al. Developing a new definition and assessing new clinical criteria for septic shock: for the Third International Consensus Definitions for Sepsis and Septic Shock (Sepsis-3). *Jama*. 2016; 315(8):775–787. <https://doi.org/10.1001/jama.2016.0289> PMID: 26903336
10. Islam MM, Nasrin T, Walther BA, Wu CC, Yang HC, Li YC. Prediction of sepsis patients using machine learning approach: a meta-analysis. *Computer methods and programs in biomedicine*. 2019; 170:1–9. <https://doi.org/10.1016/j.cmpb.2018.12.027>

11. Henry KE, Hager DN, Pronovost PJ, Saria S. A targeted real-time early warning score (TREWScore) for septic shock. *Science translational medicine*. 2015; 7(299):299ra122–299ra122. <https://doi.org/10.1126/scitranslmed.aab3719>
12. Ghosh S, Li J, Cao L, Ramamohanarao K. Septic shock prediction for ICU patients via coupled HMM walking on sequential contrast patterns. *Journal of biomedical informatics*. 2017; 66:19–31. <https://doi.org/10.1016/j.jbi.2016.12.010>
13. Calvert JS, Price DA, Chettipally UK, Barton CW, Feldman MD, Hoffman JL, et al. A computational approach to early sepsis detection. *Computers in biology and medicine*. 2016; 74:69–73. <https://doi.org/10.1016/j.compbiomed.2016.05.003> PMID: 27208704
14. Desautels T, Calvert J, Hoffman J, Jay M, Kerem Y, Shieh L, et al. Prediction of sepsis in the intensive care unit with minimal electronic health record data: a machine learning approach. *JMIR medical informatics*. 2016; 4(3). <https://doi.org/10.2196/medinform.5909> PMID: 27694098
15. Shimabukuro DW, Barton CW, Feldman MD, Mataraso SJ, Das R. Effect of a machine learning-based severe sepsis prediction algorithm on patient survival and hospital length of stay: a randomised clinical trial. *BMJ open respiratory research*. 2017; 4(1):e000234. <https://doi.org/10.1136/bmjresp-2017-000234>
16. Mao Q, Jay M, Hoffman JL, Calvert J, Barton C, Shimabukuro D, et al. Multicentre validation of a sepsis prediction algorithm using only vital sign data in the emergency department, general ward and ICU. *BMJ open*. 2018; 8(1):e017833. <https://doi.org/10.1136/bmjopen-2017-017833> PMID: 29374661
17. Delahanty RJ, Alvarez J, Flynn LM, Sherwin RL, Jones SS. Development and evaluation of a machine learning model for the early identification of patients at risk for sepsis. *Annals of emergency medicine*. 2019; 73(4):334–344. <https://doi.org/10.1016/j.annemergmed.2018.11.036>
18. Futoma J, Hariharan S, Sendak M, Brajer N, Clement M, Bedoya A, et al. An improved multi-output gaussian process rnn with real-time validation for early sepsis detection. *arXiv preprint arXiv:170805894*. 2017;.
19. Raghu A, Komorowski M, Singh S. Model-based reinforcement learning for sepsis treatment. *arXiv preprint arXiv:181109602*. 2018;.
20. Kam HJ, Kim HY. Learning representations for the early detection of sepsis with deep neural networks. *Computers in biology and medicine*. 2017; 89:248–255. <https://doi.org/10.1016/j.compbiomed.2017.08.015>
21. Moor M, Horn M, Rieck B, Roqueiro D, Borgwardt K. Early Recognition of Sepsis with Gaussian Process Temporal Convolutional Networks and Dynamic Time Warping. *arXiv preprint arXiv:190201659*. 2019;.
22. Lipton ZC, Kale DC, Wetzel R. Modeling missing data in clinical time series with rnns. *arXiv preprint arXiv:160604130*. 2016;.
23. Ghassemi M, Pimentel MA, Naumann T, Brennan T, Clifton DA, Szolovits P, et al. A multivariate time-series modeling approach to severity of illness assessment and forecasting in ICU with sparse, heterogeneous clinical data. In: *Twenty-Ninth AAAI Conference on Artificial Intelligence*; 2015.
24. Futoma J, Hariharan S, Heller K. Learning to detect sepsis with a multitask Gaussian process RNN classifier. In: *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org; 2017. p. 1174–1182.
25. Bahdanau D, Cho K, Bengio Y. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:14090473*. 2014;.
26. Yang Z, Yang D, Dyer C, He X, Smola A, Hovy E. Hierarchical attention networks for document classification. In: *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: human language technologies*; 2016. p. 1480–1489.
27. Yu L, Lin Z, Shen X, Yang J, Lu X, Bansal M, et al. MAttNet: Modular attention network for referring expression comprehension. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*; 2018. p. 1307–1315.
28. Mnih V, Heess N, Graves A, et al. Recurrent models of visual attention. In: *Advances in neural information processing systems*; 2014. p. 2204–2212.
29. Schlemper J, Oktay O, Schaap M, Heinrich M, Kainz B, Glocker B, et al. Attention gated networks: Learning to leverage salient regions in medical images. *Medical image analysis*. 2019; 53:197–207. <https://doi.org/10.1016/j.media.2019.01.012> PMID: 30802813
30. Qin Y, Song D, Chen H, Cheng W, Jiang G, Cottrell G. A dual-stage attention-based recurrent neural network for time series prediction. *arXiv preprint arXiv:170402971*. 2017;.
31. Choi E, Bahadori MT, Sun J, Kulas J, Schuetz A, Stewart W. Retain: An interpretable predictive model for healthcare using reverse time attention mechanism. In: *Advances in Neural Information Processing Systems*; 2016. p. 3504–3512.

32. Lin L, Xu B, Wu W, Richardson T, Bernal EA. Medical Time Series Classification with Hierarchical Attention-based Temporal Convolutional Networks: A Case Study of Myotonic Dystrophy Diagnosis. *arXiv preprint arXiv:190311748*. 2019;1.
33. Kong G, Lin K, Hu Y. Using machine learning methods to predict in-hospital mortality of sepsis patients in the ICU. *BMC Medical Informatics and Decision Making*. 2020; 20(1):1–10.
34. Hou N, Li M, He L, Xie B, Wang L, Zhang R, et al. Predicting 30-days mortality for MIMIC-III patients with sepsis-3: a machine learning approach using XGboost. *Journal of Translational Medicine*. 2020; 18(1):1–14. <https://doi.org/10.1186/s12967-020-02620-5> PMID: 33287854
35. Aşuroğlu T, Oğul H. A deep learning approach for sepsis monitoring via severity score estimation. *Computer Methods and Programs in Biomedicine*. 2021; 198:105816. <https://doi.org/10.1016/j.cmpb.2020.105816>
36. Yao Rq, Jin X, Wang Gw, Yu Y, Wu Gs, Zhu Yb, et al. A machine learning-based prediction of hospital mortality in patients with postoperative sepsis. *Frontiers in Medicine*. 2020; 7:445. <https://doi.org/10.3389/fmed.2020.00445>
37. Kok C, Jahmunah V, Oh SL, Zhou X, Gururajan R, Tao X, et al. Automated prediction of sepsis using temporal convolutional network. *Computers in Biology and Medicine*. 2020; 127:103957. <https://doi.org/10.1016/j.compbiomed.2020.103957> PMID: 32938540
38. Li Q, Li L, Zhong J, Huang LF. Real-time sepsis severity prediction on knowledge graph deep learning networks for the intensive care unit. *Journal of Visual Communication and Image Representation*. 2020; 72:102901. <https://doi.org/10.1016/j.jvcir.2020.102901>
39. Song W, Jung SY, Baek H, Choi CW, Jung YH, Yoo S. A Predictive Model Based on Machine Learning for the Early Detection of Late-Onset Neonatal Sepsis: Development and Observational Study. *JMIR Medical Informatics*. 2020; 8(7):e15965. <https://doi.org/10.2196/15965>
40. Svenson P, Haralabopoulos G, Torres MT. Sepsis Deterioration Prediction Using Channelled Long Short-Term Memory Networks. In: *International Conference on Artificial Intelligence in Medicine*. Springer; 2020. p. 359–370.
41. Lauritsen SM, Kalør ME, Kongsgaard EL, Lauritsen KM, Jørgensen MJ, Lange J, et al. Early detection of sepsis utilizing deep learning on electronic health record event sequences. *Artificial Intelligence in Medicine*. 2020; 104:101820. <https://doi.org/10.1016/j.artmed.2020.101820> PMID: 32498999
42. Narayanaswamy L, Garg D, Narra B, Narayanswamy R. Machine Learning Algorithmic and System Level Considerations for Early Prediction of Sepsis. In: *2019 Computing in Cardiology (CinC)*. IEEE; 2019. p. Page–1.
43. Chaudhary P, Gupta DK, Singh S. Outcome Prediction of Patients for Different Stages of Sepsis Using Machine Learning Models. In: *Advances in Communication and Computational Technology*. Springer; 2021. p. 1085–1098.
44. Bonilla EV, Chai KM, Williams C. Multi-task Gaussian process prediction. In: *Advances in neural information processing systems*; 2008. p. 153–160.
45. Li SCX, Marlin BM. A scalable end-to-end gaussian process adapter for irregularly sampled time series classification. In: *Advances in neural information processing systems*; 2016. p. 1804–1812.
46. Bai S, Kolter JZ, Koltun V. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:180301271*. 2018;.
47. Lea C, Flynn MD, Vidal R, Reiter A, Hager GD. Temporal convolutional networks for action segmentation and detection. In: *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*; 2017. p. 156–165.
48. Johnson AE, Pollard TJ, Shen L, Li-wei HL, Feng M, Ghassemi M, et al. MIMIC-III, a freely accessible critical care database. *Scientific data*. 2016; 3:160035. <https://doi.org/10.1038/sdata.2016.35> PMID: 27219127
49. Johnson A, Pollard T. sepsis3-mimic; 2018. Available from: <https://doi.org/10.5281/zenodo.1256723>.
50. Calvert J, Hoffman J, Barton C, Shimabukuro D, Ries M, Chettipally U, et al. Cost and mortality impact of an algorithm-driven sepsis prediction system. *Journal of medical economics*. 2017; 20(6):646–651. <https://doi.org/10.1080/13696998.2017.1307203> PMID: 28294646
51. Morelli A, D'Egidio A, Passariello M. Tachycardia in septic shock: pathophysiological implications and pharmacological treatment. In: *Annual Update in Intensive Care and Emergency Medicine 2015*. Springer; 2015. p. 115–128.
52. Walborn A, Williams M, Fareed J, Hoppensteadt D. International normalized ratio relevance to the observed coagulation abnormalities in warfarin treatment and disseminated intravascular coagulation. *Clinical and Applied Thrombosis/Hemostasis*. 2018; 24(7):1033–1041. <https://doi.org/10.1177/1076029618772353>

53. Fortuin V, Strathmann H, Rätsch G. Meta-Learning Mean Functions for Gaussian Processes. arXiv e-prints. 2019; p. arXiv–1901.
54. Rothfuss J, Fortuin V, Krause A. PACOH: Bayes-Optimal Meta-Learning with PAC-Guarantees. arXiv preprint arXiv:200205551. 2020;.
55. Fortuin V, Rätsch G, Mandt S. Multivariate Time Series Imputation with Variational Autoencoders. arXiv preprint arXiv:190704155. 2019;.
56. Jazbec M, Fortuin V, Pearce M, Mandt S, Rätsch G. Scalable gaussian process variational autoencoders. arXiv preprint arXiv:201013472. 2020;.
57. Ashman M, So J, Tebbutt W, Fortuin V, Pearce M, Turner RE. Sparse Gaussian Process Variational Autoencoders. arXiv preprint arXiv:201010177. 2020;.
58. Jazbec M, Pearce M, Fortuin V. Factorized Gaussian Process Variational Autoencoders. arXiv preprint arXiv:201107255. 2020;.
59. Bing S, Fortuin V, Rätsch G. On Disentanglement in Gaussian Process Variational Autoencoders. arXiv preprint arXiv:210205507. 2021;.
60. Fortuin V, Hüser M, Locatello F, Strathmann H, Rätsch G. SOM-VAE: Interpretable Discrete Representation Learning on Time Series. arXiv preprint arXiv:180602199. 2018;.
61. Manduchi L, Hüser M, Vogt J, Rätsch G, Fortuin V. DPSOM: Deep probabilistic clustering with self-organizing maps. arXiv preprint arXiv:191001590. 2019;.
62. Ciosek K, Fortuin V, Tomioka R, Hofmann K, Turner R. Conservative uncertainty estimation by fitting prior networks. In: International Conference on Learning Representations; 2019.
63. Fortuin V, Garriga-Alonso A, Wenzel F, Rätsch G, Turner R, van der Wilk M, et al. Bayesian Neural Network Priors Revisited. arXiv preprint arXiv:210206571. 2021;.
64. Garriga-Alonso A, Fortuin V. Exact Langevin Dynamics with Stochastic Gradients. arXiv preprint arXiv:210201691. 2021;.