

Early Recognition of Sepsis with Gaussian Process Temporal Convolutional Networks and Dynamic Time Warping

Michael Moor^{*,†}

MICHAEL.MOOR@BSSE.ETHZ.CH

Max Horn^{*,†}

MAX.HORN@BSSE.ETHZ.CH

Bastian Rieck^{*,†}

BASTIAN.RIECK@BSSE.ETHZ.CH

Damian Roqueiro^{*,†}

DAMIAN.ROQUEIRO@BSSE.ETHZ.CH

Karsten Borgwardt^{*,†}

KARSTEN.BORGWARDT@BSSE.ETHZ.CH

**Department of Biosystems Science and Engineering, ETH Zurich, Switzerland*

†SIB Swiss Institute of Bioinformatics, Switzerland

Abstract

Sepsis is a life-threatening host response to infection that is associated with high mortality, morbidity, and health costs. Its management is highly time-sensitive because each hour of delayed treatment increases mortality due to irreversible organ damage. Meanwhile, despite decades of clinical research, robust biomarkers for sepsis are missing. Therefore, detecting sepsis early by utilizing the affluence of high-resolution intensive care records has become a challenging machine learning problem. Recent advances in deep learning and data mining promise to deliver a powerful set of tools to efficiently address this task. This empirical study proposes two novel approaches for the early detection of sepsis: a deep learning model and a lazy learner that is based on time series distances. Our deep learning model employs a temporal convolutional network that is embedded in a multi-task Gaussian Process adapter framework, making it directly applicable to irregularly-spaced time series data. In contrast, our lazy learner is an ensemble approach that employs dynamic time warping. We frame the timely detection of sepsis as a supervised time series classification task. Consequently, we derive the most recent sepsis definition in an hourly resolution to provide the first fully accessible early sepsis detection environment. Seven hours before sepsis onset, our methods improve area under the precision–recall curve from 0.25 to 0.35 and 0.40, respectively, over the state of the art. This demonstrates that they are well-suited for detecting sepsis in the crucial earlier stages when management is most effective.

1. Introduction

Sepsis is defined as a life-threatening organ dysfunction that is caused by a dysregulated host response to infection (Singer *et al.*, 2016). Despite decades of clinical research, sepsis remains a major public health issue that is associated with high mortality, morbidity, and related health costs (Dellinger *et al.*, 2013; Kaukonen *et al.*, 2014; Hotchkiss *et al.*, 2016).

Currently, when sepsis is detected and the underlying pathogen is identified, organ damage has already progressed to a potentially irreversible stage. Effective management, especially in the intensive care unit (ICU), is of critical importance. From sepsis onset, each hour of delayed effective antibiotic treatment increases mortality (Ferrer *et al.*, 2014). Therefore, early detection of sepsis has gained considerable attention in the machine learning community (Kam and Kim, 2017; Futoma *et al.*, 2017a). The task of detecting sepsis

early has often been modeled as a multi-channel time series classification task. Clinical data is commonly sampled irregularly, thus often requiring a set of hand-crafted preprocessing steps, such as binning, carry-forward imputation, and rolling means (Calvert *et al.*, 2016; Desautels *et al.*, 2016) prior to the application of a predictive model. However, these imputation schemes lead to a loss of data sparsity, which may carry crucial information in this context. Most existing approaches are incapable of retaining sampling information, thereby potentially impeding the training and leading to lower predictive performance. Futoma *et al.* (2017a) proposed a sepsis detection method that accounts for irregular sampling by applying the Gaussian Process adapter end-to-end learning framework (Li and Marlin, 2016) and then training it using a long short-term memory (LSTM) classifier (Hochreiter and Schmidhuber, 1997). Only recently have convolutional networks gained attention in sequence modeling (Gehring *et al.*, 2017; Vaswani *et al.*, 2017). In particular, temporal convolutional networks (TCNs; Lea *et al.* (2017)), have been shown to outperform conventional recurrent neural network (RNN) architectures for many sequential learning tasks in terms of evaluation metrics, memory efficiency, and parallelism (Bai *et al.*, 2018). In light of these developments, we propose a deep learning model as an end-to-end trainable framework for early sepsis detection that builds on both multi-task Gaussian Process (MGP) adapters (which are an extension of Gaussian Process adapters to multi-task learning) and TCNs. We refer to this model as MGP-TCN because it combines the uncertainty-aware framework of GP adapters with TCNs. The contributions of our work are threefold:

- We present a lazy learner that is based on dynamic time warping and k -nearest neighbors (DTW-KNN), which can be seen as a multi-channel ensemble extension of a well-established data mining technique for time series classification. Moreover, we develop MGP-TCN, which is the first model that can leverage temporal convolutions on irregularly-sampled multivariate time series.
- We provide the first fully-accessible framework for the early detection of sepsis on a benchmark dataset featuring a publicly available temporally resolved Sepsis-3 label to enable community-based sepsis detection research¹.
- We present a detailed experimental setup in which we empirically demonstrate that our methods outperform the state of the art in detecting sepsis *early*.

Technical Significance Ours is the first work to combine the MGP adapter framework (Bonilla *et al.*, 2008; Li and Marlin, 2016) with TCNs (Lea *et al.*, 2017), thus improving memory efficiency, scalability, and classification performance for sepsis early detection on irregularly-sampled time series. We outperform the state-of-the-art method and improve AUPRC from 0.25 to 0.35/0.40, respectively (measured 7 h before sepsis onset).

Clinical Relevance The delayed identification and treatment of sepsis is a major driver of mortality in the ICU. By detecting sepsis *earlier*, our approach could significantly decrease mortality, because timely management is essential in this context (Ferrer *et al.*, 2014). An early warning system that is based on our methods could prevent delays in the initiation of antimicrobial and supportive therapy, which are considered to be crucial for improving patient outcome (Kumar *et al.*, 2006).

1. See <https://github.com/BorgwardtLab/mgp-tcn> for more details.

2. Related Work

This section introduces the recent literature and current challenges for sepsis detection.

2.1. Supervised Learning on Medical Time Series

Supervised learning on time series datasets has been haunted by the crux that labels per time point are often missing, especially in medical applications (Reddy and Aggarwal, 2015). This hindrance also applies to the early detection of sepsis. In previous work, it was usually circumvented by applying ad-hoc schemes to determine resolved sepsis labels (Calvert *et al.*, 2016; Mao *et al.*, 2018; Kam and Kim, 2017). These papers used a global time series label, such as an ICD disease code intended for billing, and they estimated sepsis onset with easily computable ad-hoc criteria. However, when using such a patchwork label, it is unclear if the patient *actually* suffered from an event at this time and not, for instance, one week later.

By extracting Sepsis-3, which is the most recent sepsis criterion (Singer *et al.*, 2016) that allows for temporal resolution, we contribute a solution to this issue that continues to affect the study of machine learning for healthcare. Even though some datasets (Futoma *et al.*, 2017a; Desautels *et al.*, 2016) have high-resolution sepsis labels, they are currently *not* accessible to the research community. This leads to reproducibility and comparability issues. Thus, there are massive hurdles to overcome before novel approaches for sepsis detection can be developed and thoroughly validated.

2.2. Algorithms for the Early Detection of Sepsis

Overview In the last decade, several data-driven approaches for detecting sepsis in the ICU have been presented (Desautels *et al.*, 2016; Calvert *et al.*, 2016; Kam and Kim, 2017; Futoma *et al.*, 2017a; Shashikumar *et al.*, 2017). Many approaches selectively compare with simple clinical scores, such as SIRS, NEWS or MEWS (Bone *et al.*, 1992; Williams *et al.*, 2012; Stenhouse *et al.*, 2000). However, none of these scores are intended as specific, continuously-evaluated risk scores for sepsis. Specifically, the SIRS criteria are now considered by clinicians to be unspecific and obsolete for the definition of sepsis (Beesley and Lanspa, 2015; Kaukonen *et al.*, 2015). As an alternative to these scores, Henry *et al.* (2015) introduced a targeted real-time warning score (TREWScore) to predict septic shock, which is a frequent complication following from sepsis. Notably, while many machine learning methods have surpassed generic or simplistic clinical schemes, next to no papers actually performed the hard comparison to other machine learning approaches in the literature. As an exception, the application of LSTMs (Kam and Kim, 2017) have been shown to be an improvement over the InSight model (Calvert *et al.*, 2016), which is a regression model with hand-crafted features. However, only reported metrics were compared, whereas potentially differing processing pipelines and label implementations (which are closed source) could make a direct comparison problematic. These circumstances prompted us to baseline our work against state-of-the-art machine learning methods and on exactly the same sepsis early recognition pipeline, which we make publicly available.

State of the Art Sepsis detection methods are usually developed on real-world datasets with prevalence values ranging from 6.6% (Kam and Kim, 2017) to 21.4% (Futoma *et al.*, 2017a). Despite this considerable class imbalance, to our knowledge, only Futoma *et al.*

(2017a) and Desautels *et al.* (2016) report the area under the precision–recall curve (AUPRC), in addition to the area under the receiver operating characteristic curve (AUC). Given the class imbalance, AUC is known to be a less informative evaluation criterion (Saito and Rehmsmeier, 2015). Thus, in terms of AUPRC, Futoma *et al.* (2017a) currently represent the state of the art in the early detection of sepsis. In a follow-up paper, Futoma *et al.* (2017b) improved their performance by proposing task-specific tweaks, such as label-propagation, additional feature extraction (e.g. missingness indicators), and separate task correlation matrices for their Gaussian Process. However, these extensions pertain to the input features and they modify the GP adapter framework as wrapped around their classifier, so they are orthogonal to our undertaking of improving the *classifier* inside the GP adapter framework.

2.3. Gaussian Process Adapters

Li and Marlin (2016) showed that optimizing a Gaussian Process imputation of a time series end-to-end using the gradients of a subsequent classifier leads to better performance than optimizing both the classifier and the GP separately. This method, which is also referred to as GP adapters, is not restricted to imputing missing data (Li *et al.*, 2017). Recently, Futoma *et al.* (2017a) demonstrated that GP adapters are a well-suited framework to handle the irregularly spaced time series in early sepsis detection. Specifically, they confirmed earlier findings (Li and Marlin, 2016) that in time series classification, GP adapters outperform conventional GP imputation schemes that require a separate optimization step, which is not driven by the classification task.

3. Methods

In the following, we describe our proposed MGP-TCN and DTW-KNN methods². First, Section 3.1 gives a high-level overview of our deep learning method³, emphasizing the MGP component (i.e., the first building block of the method) which as a whole was previously applied by Futoma *et al.* (2017a). Section 3.2 then describes temporal convolutional networks (TCNs), the second building block. Finally, Section 3.3 describes DTW-KNN.

3.1. Multi-task Gaussian Process Temporal Convolutional Network Classifier

We frame the early detection of sepsis in the ICU as a multivariate time series classification task. Specifically, we focus on the task of identifying sepsis onset in irregularly-sampled multivariate time series of physiological measurements in ICU patients. Our proposed model uses a multi-task Gaussian Process (MGP) (Bonilla *et al.*, 2008) that is intrinsically capable of dealing with non-uniform sampling frequencies. In this setting, the tasks considered by the MGP are the individual channels of the time series. More precisely, given irregularly-observed measurements (values and times) $\{\mathbf{y}_i, \mathbf{t}_i\}$ of encounter i , for evenly-spaced query times \mathbf{x}_i , the MGP draws a latent time series \mathbf{z}_i following the MGP’s posterior distribution $\mathcal{P}(\mathbf{z}_i | \mathbf{y}_i, \mathbf{t}_i, \mathbf{x}_i; \boldsymbol{\theta})$ (see Equation 1). \mathbf{z}_i then serves as the input to a temporal convolutional network (TCN, Section 3.2) that predicts the sepsis label. Making use of the Gaussian

2. Our notation uses regular font for scalars, bold lower-case for vectors, and bold upper-case for matrices.

3. Please refer to Supplementary Section A.1 for more details on the end-to-end MGP adapter framework.




Figure 1: Overview of our model. Raw, irregularly-spaced time series are provided to the multi-task Gaussian Process (MGP) for each patient. The MGP then draws from a posterior distribution, given the observed data, at evenly-spaced grid times (each hour). This grid is then fed into a temporal convolutional network (TCN) which, after a forward pass, returns a loss. Its gradient is then computed by backpropagation through both the TCN and the MGP (green arrows). All parameters are learned end-to-end during training.

Process adapter framework (Li and Marlin, 2016) enables us to optimize this entire process end-to-end with respect to the final classification objective; that is, identifying sepsis. Figure 1 gives a high-level overview of the MGP-TCN model. The MGP’s posterior distribution follows a multivariate normal distribution, i.e.

$$\mathbf{z}_i \sim \mathcal{N}(\boldsymbol{\mu}(\mathbf{z}_i), \boldsymbol{\Sigma}(\mathbf{z}_i); \boldsymbol{\theta}), \quad (1)$$

with mean and covariance

$$\boldsymbol{\mu}(\mathbf{z}_i) = (\mathbf{K}^D \otimes \mathbf{K}^{X_i T_i})(\mathbf{K}^D \otimes \mathbf{K}^{T_i} + \mathbf{D} \otimes \mathbf{I})^{-1} \mathbf{y}_i \quad (2)$$

$$\boldsymbol{\Sigma}(\mathbf{z}_i) = (\mathbf{K}^D \otimes \mathbf{K}^{X_i}) - (\mathbf{K}^D \otimes \mathbf{K}^{X_i T_i})(\mathbf{K}^D \otimes \mathbf{K}^{T_i} + \mathbf{D} \otimes \mathbf{I})^{-1}(\mathbf{K}^D \otimes \mathbf{K}^{T_i X_i}). \quad (3)$$

Here, $\mathbf{K}^{X_i T_i}$ refers to the correlation matrix between the evenly-spaced query times \mathbf{x}_i and the observed times \mathbf{t}_i , while \mathbf{K}^{X_i} represents the correlations between \mathbf{x}_i with itself. \mathbf{K}^D is the task-similarity kernel matrix whose entry $K_{d,d'}^D$ at position (d, d') represents the similarity of tasks (i.e., time series channels) d and d' . \otimes denotes the Kronecker product, and \mathbf{K}^{T_i} represents an encounter-specific $T_i \times T_i$ correlation matrix between all observed times $t_i \in \mathbf{t}_i$ of patient encounter i , while \mathbf{D} is a diagonal matrix of per-task noise variances satisfying $D_{dd} = \sigma_d^2$ and \mathbf{I} refers to the identity matrix. The posterior mean $\boldsymbol{\mu}(\mathbf{z}_i)$ also depends on the observed values \mathbf{y}_i . We gather the MGP’s parameters in $\boldsymbol{\theta} = \{\mathbf{K}^D, \sigma_d^2 |_{d=1}^D, l\}$ where l refers to the length scale of the kernel function. For more details, please refer to Section A.1.

3.2. Temporal Convolutional Networks

This section outlines the details of a generic temporal convolutional network (TCN) architecture. TCNs have recently been proposed (Lea et al., 2017) as an extension of convolutional neural networks (CNNs), which are known to exhibit state-of-the-art performance in visual




Figure 2: Schematic illustration of the TCN architecture. The input \mathbf{z}_i values (blue) of the TCN classifier are computed by the multi-task Gaussian Process on a regular grid (x_0, \dots, x_t) based on the observed values (yellow). Each temporal block skips an increasing number of the previous layer’s outputs, such that the *visible window* of a single node increases exponentially with increasing number of layers. Figure recreated from Bai *et al.* (2018).

tasks (Cireřan *et al.*, 2011, 2012). An empirical study by Bai *et al.* (2018) demonstrated that TCNs show superior performance for sequence modeling tasks, as compared to recurrent neural networks. Please see Figure 2 for an illustration of our TCN architecture, for which the subsequent sections provide more details.

Causal Dilated Convolutions TCNs are a simple but powerful extension to conventional 1D-CNNs in that they exhibit three properties (Bai *et al.*, 2018):

1. **Sequence to sequence:** The output of a TCN has the same length as its input.
2. **Convolutions are causal:** Outputs are only influenced by present and past inputs.
3. **Long effective memory:** By using dilated convolutions, the receptive field, and thus also the effective memory, grows exponentially with increasing network depth.

For 1. and 2., we use zero-padding to enforce equal layer sizes throughout all layers while ensuring that for the output at time t , only input values at this time and earlier times can be used (see Figure 2). For 3., we follow the approach of Yu and Koltun (2015) by defining the l -dilated convolution of an input sequence \mathbf{x} with a filter \mathbf{f} as

$$(\mathbf{x} *_l \mathbf{f})(k) = \sum_{j=i+l} x_i \cdot f_j, \tag{4}$$

where the 1-dilated convolution coincides with the regular convolution. By stacking l -dilated convolutions in an exponential manner, such that $l = 2^n$ for the n -th layer, a long effective memory can be achieved, as illustrated by Bai *et al.* (2018).




Figure 3: For each encounter with a suspicion of infection (SI), we extract a 72 h window around the first SI event (starting 48 h before) as the SI-window. The Sequential Organ Failure Assessment (SOFA) score is then evaluated for every hour in this window by combining physiological scores of six organ systems. Following the SOFA definition, to arrive at a SOFA score we considered the worst organ scores of the last 24 h.

Residual Temporal Blocks We stack convolutional layers of a TCN into residual temporal blocks; that is, blocks that combine the previous input and the result of the respective convolution with an addition. Thus, the output of a temporal block is computed relatively with respect to an input. Here, we follow the setup of Lee (2018), which applies layer normalization (Ba *et al.*, 2016) to improve training stability and convergence, as opposed to the weight normalization employed by Bai *et al.* (2018). Furthermore, we apply normalization *after* each activation, similar to Lee (2018).

3.3. Dynamic Time Warping Classifier

Dynamic time warping (Keogh and Pazzani, 1999) for time series classification, using k -nearest neighbor approaches (here referred to as DTW-KNN), is known to exhibit highly competitive predictive performance (Dau *et al.*, 2017; Ding *et al.*, 2008). As opposed to many other off-the-shelf classifiers, it can handle variable-length time series. Despite its wide-spread use and demonstrated capabilities in data mining, DTW-KNN has (to the best of our knowledge) never been used in sepsis detection tasks. We thus extend DTW-KNN for the classification of multivariate time series, thereby introducing an additional novel approach for the early detection of sepsis. More precisely, we address the multivariate nature of our setup by computing the DTW distance matrix (containing the pairwise distances between all patients) for each time series channel separately. Each distance matrix is subsequently used for training a k -nearest neighbor classifier. Instead of using resampling techniques, the ensemble is constructed by combining *all* per-channel classifiers. Finally, for the classification step, the final prediction score is computed as the average over all per-channel prediction scores.

4. Experiments

4.1. Dataset and Sepsis Label Definition

Our analysis uses the MIMIC-III (*Multiparameter Intelligent Monitoring in Intensive Care*) database, version 1.4 (Johnson *et al.*, 2016). MIMIC-III includes over 58,000 hospital admissions of over 45,000 patients, as encountered between June 2001 and October 2012. We follow the most recent sepsis definition (Singer *et al.*, 2016), which requires a co-occurrence of suspected infection (SI) and organ dysfunction. For SI, we follow the recommendations of Seymour *et al.* (2016) to implement the SI cohort (please refer to Supplementary Section A.2.2 for more details).

According to Singer *et al.* (2016), the organ dysfunction criterion is fulfilled when the SOFA Score (Vincent *et al.*, 1996) shows an increase of at least 2 points. To determine this, we follow the suggestions of Singer *et al.* to use a window of -48 h to 24 h around a suspicion of infection. Figure 3 illustrates our Sepsis-3 implementation. To detect sepsis *early*, determining the sepsis onset time is crucial. We thus considerably refined and extended the queries provided by Johnson *et al.* (2018) to determine the Sepsis-3 label on an hourly basis⁴. If sepsis is determined by merely checking whether a patient fulfills the criteria upon admission, similarly to how it is done by Johnson *et al.* (2018), then only those patients who *arrive* in the ICU with sepsis would be defined as cases, not the—arguably more interesting ones—that *develop* the syndrome during their ICU stay.

4.2. Data Filtering

Patient Inclusion Criteria We exclude patients under the age of 15 and those for which no chart data is available—including ICU admission or discharge time. Furthermore, following the recent sepsis literature, ICU encounters logged via the CareVue system were excluded due to underreported negative lab measurements (Desautels *et al.*, 2016). We include an encounter as a case if at any time during the ICU stay a sepsis onset occurs, whereas controls are defined as those patients that have no sepsis onset (they still might have suspected infection or organ dysfunction, separately). To additionally ensure that controls cannot be sepsis patients that developed sepsis shortly before ICU, we require controls not to be labeled with any sepsis-related ICD-9 billing code. Following these inclusion criteria, we initially count 1,797 sepsis cases and 17,276 controls. This work aims for sepsis *early* detection, so we follow Desautels *et al.* (2016) and exclude cases that develop sepsis earlier than seven hours into the ICU stay. This enables a prediction horizon of 7 h. To preserve a realistic class balance of around 10%, we apply this exclusion step only after the case-control matching (see next paragraph). Thus, after cleaning and filtering, we finally use 570 cases and 5,618 controls as our cohort; Table 1 shows the summary statistics. For the variables, we used 44 irregularly-sampled laboratory and vital parameters⁵. Furthermore, to be able to run all baselines, we had to apply an additional patient filtering step⁶.

4. Their provided code only checks whether a simplified version of Sepsis-3 is satisfied upon admission. For instance, no *increase* in SOFA points is considered, but only one abnormally high value.

5. Please see Supplementary Section A.3 for more details.

6. Please see Supplementary Section A.8 for more details.

Table 1: Characteristics of the population included in the dataset. The mean sepsis onset is given in hours since admission to the ICU.

Variable	Sepsis Cases	Controls
n	570	5,618
Female	236 (41.4%)	2,548 (45.4%)
Male	334 (58.6%)	3,070 (54.6%)
Mean time to sepsis onset in ICU (median)	16.7 h (11.8 h)	—
Age ($\mu \pm \sigma$)	67.2 \pm 15.3	64.2 \pm 17.3
Ethnicity		
White	411 (72.1%)	4,047 (72.0%)
Black or African-American	41 (7.2%)	551 (9.8%)
Hispanic or Latino	7 (1.2%)	147 (2.6%)
Other	57 (10.0%)	493 (8.8%)
Not available	54 (9.5%)	380 (6.8%)
Admission type		
Emergency	504 (88.4%)	4,689 (83.5%)
Elective	60 (10.5%)	872 (15.5%)
Urgent	6 (1.1%)	57 (1.0%)

Case-control Matching In previous work, it has been observed that an insufficient alignment of time series of sepsis cases versus controls could render the classification task trivial: for instance, when comparing a window before sepsis onset to the last window (before discharge) of a control’s ICU stay, the classification task is much easier than when compared to a more similar reference time in a control’s stay. This can be observed by the decrease in performance of the MGP-RNN method when Futoma *et al.* (2017b) applied case-control matching. Hence, to avoid a trivial classification task, we also use a case-control alignment in a matching procedure. To accommodate the class imbalance, we assign each case to 10 random unassigned controls and define their control onset as the absolute time (in hours since admission) when the matched case fulfilled the sepsis criteria. Futoma *et al.* (2017b) used a relative measure; that is, the same percentage of the entire ICU stay as for control onset time. However, we observed that cases and controls do not necessarily share the same length of stay, which could introduce bias to the alignment that a deep neural network could potentially exploit. For each case and its matched controls, we extract up to 48 h of input data preceding their onset and after ICU admission.

4.3. Experimental Setup

Baselines We compare our methods against MGP-RNN, which is the current state-of-the-art sepsis detection method (Futoma *et al.*, 2017a,b). To enable a fair comparison, the authors kindly provided source code for their pipeline such that their model could be trained

from scratch on our dataset. Additionally, we compare against a classical TCN (here referred to as Raw-TCN) that is not embedded in the MGP adapter framework. To this end, as a preprocessing step, we first impute the times series using a carry-forward scheme (for more details, please refer to Supplementary Section A.4). We train our DTW-KNN ensemble classifier using the same imputation scheme.

Training We apply three iterations of random splitting using 80% of the samples for training and each 10% for validation and testing. In each random split, the time series were z -scored per channel using the respective train mean and standard deviation. For hyperparameter tuning, due to the costly evaluations, we apply an off-the-shelf Bayesian optimization framework (instead of an exhaustive grid search) provided by the `scikit-optimize` library (`scikit-optimize contributors, 2018`) with 20 calls per method and split. We select the best model parameters and checkpoints in terms of validation AUPRC and we evaluate them on the test splits. For all deep models, the hyperparameter spaces were constrained such that the number of parameters each ranged from 20K–500K. To make it feasible to analyze multiple random splits (despite a deep learning setup), we constrain each run to take at most two hours. To prevent underfitting, we additionally retrained the best parameter setting of each method and split for a longer period of 50 epochs. Please refer to Supplementary Section A.7 for more details.

Evaluation Due to substantial class imbalance (the overall case prevalence is 9.2%, or roughly 1 case for 10 controls), we evaluate all models in terms of AUPRC on the test split. In addition, we report AUC, mostly to comply with recent sepsis detection literature (see also Section 2.2 for a discussion of the disadvantages of this measure). Because timely identification of sepsis is of central importance, we evaluate the trained models in a horizon analysis going back up to 7 h before sepsis onset. For example, to evaluate the prediction horizon at 3 h in advance, for each encounter, the model (and imputation scheme) is only provided with input data up until that moment. To assess the predictive performance, we do *not* optimize the models to each respective horizon hour, which would result in eight distinct specialized models. Instead, per method and fold, we train one model on all of the available training data and challenge its performance by gradually restricting access to the information closest to sepsis onset.

Implementation Details Additional information about the technical details of our implementation and its runtime behavior are available in Supplementary Section A.9.

4.4. Results

Figure 4 depicts the predictive performance for the different time horizons. The x -axes indicate the prediction horizon in hours before sepsis onset. The y -axes measure AUPRC (left) and AUC (right). As previously discussed, we focus on evaluating AUPRC due to the substantial class imbalance (9.2%).

We observe that both our novel model MGP-TCN and our DTW-KNN ensemble method *consistently* outperform the current state-of-the-art early sepsis detection classifier MGP-RNN. Especially for the early detection task of more than 4 h before sepsis onset, both MGP-TCN and DTW-KNN outperform the state of the art with a significant margin, while the latter shows slightly higher performance in this regime. For horizons that are closer




Figure 4: We evaluate all methods using area under the precision–recall curve (AUPRC) and additionally display the (less informative) area under the receiver operating characteristic curve (AUC). The current state-of-the-art method, MGP-RNN, is shown in blue. The two approaches for early detection of sepsis that were introduced in this paper (i.e., MGP-TCN and DTW-KNN) are shown in pink and red, respectively. Using three random splits for all measures and methods, we show the mean (line) and standard deviation error bars (shaded area).

to sepsis onset, namely 0 h to 3 h prior to onset, *all* approaches except Raw-TCN exhibit overlapping performance in terms of their variance estimates. Interestingly, Raw-TCN does not yield competitive results for *any* setting that was considered in our experimental setup. With increasing distance to the onset, its performance starts to approximate that of the MGP-RNN classifier. Finally, approaches based on simple imputation schemes (i.e., Raw-TCN and DTW-KNN) exhibit a much flatter trend in AUPRC when approaching sepsis onset than the MGP-imputed ones.

5. Conclusion

Our proposed methods MGP-TCN and DTW-KNN exhibit favorable performances over all prediction horizons and they consistently outperform the state-of-the-art baseline classifier MGP-RNN. Compared to the classic TCN, we empirically demonstrated that TCN-based architectures—in combination with MGPs to account for uncertainty associated with irregular sampling—result in competitive predictive performance. Specifically, in terms of AUPRC, with MGP-TCN and DTW-KNN, we improve the current state of the art from 0.25 to 0.35 and 0.40, respectively, 7 h before sepsis onset. This confirms that recent advances in sequence modeling may be transferred successfully to irregularly-sampled medical time series. By contrast, the low performance of the Raw-TCN classifier suggests that a more advanced, uncertainty-aware imputation scheme is helpful when transferring “deep”

approaches to our scenario. Furthermore, we observed that simple imputation schemes lead to a smaller gain in performance when approaching sepsis onset. This could be due to the nature of the carry-forward scheme, which tends to remove relevant sampling information.

When comparing our findings with the recent literature, we observe that the low prevalence in our dataset (9.2%) makes the classification task substantially harder. For instance, in terms of AUPRC, MGP-RNN performed better on its original dataset (to which we unfortunately have no access), which has a prevalence of 21.4% (Futoma *et al.*, 2017a). In terms of prevalence and preprocessing, to our knowledge the most comparable setup would be the one by Desautels *et al.* (2016); however, we have made several requests to obtain their methods and queries, which have proven to be unsuccessful. Interestingly, their reported AUPRC dropped to roughly 0.3 already at 1 h before sepsis onset, where, for example, our proposed MGP-TCN method still achieves an AUPRC of 0.51.

One of the most surprising results is the highly-competitive performance of our DTW-KNN ensemble classifier, whose performance for earlier horizons exceeds all of the other methods, despite its conceptual simplicity. While this is highly relevant for the early detection of sepsis, the DTW-KNN classifier suffers from some practical limitations that impede online monitoring scenarios and its application to very large patient cohorts. Already for our dataset, using a standard implementation, predicting at one horizon may require hundreds of millions of pairs of univariate time series to be aligned, followed by their distance computation, and the subsequent storing of results (which potentially affects both runtime and memory). We conjecture that DTW-KNN performs so well because of the mid-range sample size, whereas deep models tend to perform best for even larger sample sizes. However, scaling DTW-KNN to cohorts of hundreds of thousands of patients currently appears to be a computational challenge. This also affects online classification in the ICU: for each new measurement of a patient, the distances of each involved channel to *all* patients in the training cohort have to be updated, and partially recomputed. Consequently, intermediate results of the distance calculations need to be stored, leading to a significant memory overhead. The problem thus remains a “hot topic” in time series analysis (Oregi *et al.*, 2017).

In contrast, MGP-TCN does not suffer from these limitations because only the network weights have to be stored and classifying a new patient is *constant* in the total number of patients. Thus, MGP-TCN can be easily applied to larger cohorts, which is likely to further increase predictive performance. Moreover, obtaining online predictions only requires very slight modifications. For more details on the scaling behavior, please refer to Supplementary Section A.6.

6. Future Work

An additional source of validation of our findings would be to test our model on more datasets. For the early detection of sepsis, this is normally not done because the derivation of properly-resolved, time-based sepsis labels requires considerable preprocessing efforts. For example, in our work, the dynamic sepsis labels first required the implementation of an entire query pipeline. Due to this bottleneck, the value of providing publicly accessible sepsis labels for further research, as we do in this work, cannot be overstated. In the future, we also would like to extend our analysis to more types of data sources arising from the

ICU. Futoma *et al.* (2017b) already employed a subset of baseline covariates, medication effects, and missingness indicator variables. However, a multitude of feature classes still remain to be explored and integrated, each posing unique challenges that will be interesting to overcome. For instance, the combination of sequential and non-sequential features has previously been handled by treating non-sequential features as sequential features (Futoma *et al.*, 2017a). We hypothesize that this could be handled more efficiently by using a more modular architecture that handles sequential and non-sequential features differently. Furthermore, we aim to obtain a better understanding of the time series features used by the model. Specifically, we are interested in assessing the *interpretability* of the learned filters of the MGP-TCN framework and then evaluate how much the activity of an individual filter contributes to a prediction. This endeavor is somewhat facilitated by our use of a convolutional architecture. The extraction of short per-channel signals could prove very relevant for supporting diagnoses made by clinical practitioners.

Funding

This work was funded in part by the SPHN/PHRT Driver Project “Personalized Swiss Sepsis Study” as well as by the Alfred Krupp Prize for Young University Teachers of the Alfred Krupp von Bohlen und Halbach-Stiftung (K.B.).

References

- Ba, J. L. *et al.* (2016). Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Bai, S. *et al.* (2018). An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271*.
- Beesley, S. J. and Lanspa, M. J. (2015). Why we need a new definition of sepsis. *Annals of Translational Medicine*, **3**(19).
- Bone, R. C. *et al.* (1992). Definitions for sepsis and organ failure and guidelines for the use of innovative therapies in sepsis. *Chest*, **101**(6), 1644–1655.
- Bonilla, E. V. *et al.* (2008). Multi-task gaussian process prediction. In *Advances in Neural Information Processing Systems*, pages 153–160.
- Calvert, J. S. *et al.* (2016). A computational approach to early sepsis detection. *Computers in Biology and Medicine*, **74**, 69–73.
- Ciresan, D. C. *et al.* (2011). Flexible, high performance convolutional neural networks for image classification. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1237–1242.
- Ciresan, D. C. *et al.* (2012). Multi-column deep neural networks for image classification. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3642–3649.

- Dau, H. A. *et al.* (2017). Judicious setting of Dynamic Time Warping’s window width allows more accurate classification of time series. In *IEEE International Conference on Big Data*, pages 917–922.
- Dellinger, R. P. *et al.* (2013). Surviving sepsis campaign: International guidelines for management of severe sepsis and septic shock 2012. *Critical Care Medicine*, **41**(2), 580–637.
- Desautels, T. *et al.* (2016). Prediction of sepsis in the intensive care unit with minimal electronic health record data: A machine learning approach. *JMIR Medical Informatics*, **4**(3), e28.
- Ding, H. *et al.* (2008). Querying and mining of time series data: experimental comparison of representations and distance measures. *Proceedings of the VLDB Endowment*, **1**(2), 1542–1552.
- Ferrer, R. *et al.* (2014). Empiric antibiotic treatment reduces mortality in severe sepsis and septic shock from the first hour: results from a guideline-based performance improvement program. *Critical Care Medicine*, **42**(8), 1749–1755.
- Futoma, J. *et al.* (2017a). Learning to detect sepsis with a multitask Gaussian process RNN classifier. In *International Conference on Machine Learning*, pages 1174–1182.
- Futoma, J. *et al.* (2017b). An improved multi-output Gaussian process RNN with real-time validation for early sepsis detection. In *Proceedings of the 2nd Machine Learning for Healthcare Conference (MLHC)*.
- Gehring, J. *et al.* (2017). Convolutional sequence to sequence learning. In *International Conference on Machine Learning*, pages 1243–1252.
- Henry, K. E. *et al.* (2015). A targeted real-time early warning score (TREWScore) for septic shock. *Science Translational Medicine*, **7**(299).
- Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, **9**(8), 1735–1780.
- Hotchkiss, R. S. *et al.* (2016). Sepsis and septic shock. *Nature Reviews Disease Primers*, **2**.
- Johnson, A. E. *et al.* (2016). MIMIC-III, a freely accessible critical care database. *Scientific Data*, **3**.
- Johnson, A. E. *et al.* (2018). The MIMIC Code Repository: enabling reproducibility in critical care research. *Journal of the American Medical Informatics Association*, **25**(1), 32–39.
- Kam, H. J. and Kim, H. Y. (2017). Learning representations for the early detection of sepsis with deep neural networks. *Computers in Biology and Medicine*, **89**, 248–255.
- Kaukonen, K.-M. *et al.* (2014). Mortality related to severe sepsis and septic shock among critically ill patients in Australia and New Zealand, 2000-2012. *Journal of the American Medical Association (JAMA)*, **311**(13), 1308–1316.

- Kaukonen, K.-M. *et al.* (2015). Systemic inflammatory response syndrome criteria in defining severe sepsis. *New England Journal of Medicine*, **372**(17), 1629–1638.
- Keogh, E. J. and Pazzani, M. J. (1999). Scaling up dynamic time warping to massive datasets. In J. M. Żytkow and J. Rauch, editors, *Principles of Data Mining and Knowledge Discovery*, pages 1–11, Heidelberg, Germany. Springer.
- Kumar, A. *et al.* (2006). Duration of hypotension before initiation of effective antimicrobial therapy is the critical determinant of survival in human septic shock. *Critical Care Medicine*, **34**(6), 1589–1596.
- Lea, C. *et al.* (2017). Temporal convolutional networks for action segmentation and detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1003–1012.
- Lee, C. (2018). Solving sequential MNIST with Temporal Convolutional Networks (TCNs). <https://colab.research.google.com/drive/11a331W7FQV1RicipfzyLq9H0SH1VSD4LE#scrollTo=b10XJFD-y-fS>.
- Li, S. C.-X. and Marlin, B. M. (2016). A scalable end-to-end Gaussian process adapter for irregularly sampled time series classification. In *Advances in Neural Information Processing Systems*, pages 1804–1812.
- Li, Y. *et al.* (2017). Targeting EEG/LFP synchrony with neural nets. In *Advances in Neural Information Processing Systems*, pages 4620–4630.
- Mao, Q. *et al.* (2018). Multicentre validation of a sepsis prediction algorithm using only vital sign data in the emergency department, general ward and icu. *BMJ open*, **8**(1), e017833.
- Oregi, I. *et al.* (2017). On-line dynamic time warping for streaming time series. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 591–605. Springer.
- Reddy, C. K. and Aggarwal, C. C. (2015). *Healthcare data analytics*. Chapman and Hall/CRC.
- Saito, T. and Rehmsmeier, M. (2015). The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. *PLoS one*, **10**(3), e0118432.
- scikit-optimize contributors, T. (2018). `scikit-optimize/scikit-optimize: v0.5.2`. <https://doi.org/10.5281/zenodo.1207017>.
- Seymour, C. W. *et al.* (2016). Assessment of clinical criteria for sepsis: For the Third International Consensus Definitions for Sepsis and Septic Shock (Sepsis-3). *Journal of the American Medical Association (JAMA)*, **315**(8), 762–774.

- Shashikumar, S. *et al.* (2017). Early sepsis detection in critical care patients using multiscale blood pressure and heart rate dynamics. *Journal of Electrocardiology*.
- Singer, M. *et al.* (2016). The Third International Consensus Definitions for Sepsis and Septic Shock (Sepsis-3). *Journal of the American Medical Association (JAMA)*, **315**(8), 801–810.
- Stenhouse, C. *et al.* (2000). Prospective evaluation of a modified early warning score to aid earlier detection of patients developing critical illness on a general surgical ward. *British Journal of Anaesthesia*, **84**(5), 663P.
- Vaswani, A. *et al.* (2017). Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Vincent, J.-L. *et al.* (1996). The SOFA (sepsis-related organ failure assessment) score to describe organ dysfunction/failure. *Intensive Care Medicine*, **22**(7), 707–710.
- Williams, B. *et al.* (2012). National Early Warning Score (NEWS): Standardising the assessment of acute-illness severity in the NHS. *London: The Royal College of Physicians*.
- Yu, F. and Koltun, V. (2015). Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*.

Table A.1: List of all 44 used clinical variables. For this study, we focused on irregularly sampled time series data comprising vital and laboratory parameters. To exclude variables that rarely occur, we selected variables with 500 or more observations present in the patients that fulfilled our original inclusion criteria (1,797 cases and 17,276 controls).

Vital Parameters	
Systolic Blood Pressure	Tidal Volume Set
Diastolic Blood Pressure	Tidal Volume Observed
Mean Blood Pressure	Tidal Volume Spontaneous
Respiratory Rate	Peak Inspiratory Pressure
Heart Rate	Total Peep Level
SpO ₂ (Pulsoxymetry)	O ₂ flow
Temperature Celsius	FiO ₂ (Fraction of Inspired Oxygen)
Cardiac Output	
Laboratory Parameters	
Albumin	Blood Urea Nitrogen
Bands (Immature Neutrophils)	White Blood Cells
Bicarbonate	Creatine Kinase
Bilirubin	Creatine Kinase MB
Creatinine	Fibrinogen
Chloride	Lactate Dehydrogenase
Sodium	Magnesium
Potassium	Calcium (free)
Lactate	pO ₂ Bloodgas
Hematocrit	pH Bloodgas
Hemoglobin	pCO ₂ Bloodgas
Platelet Count	SO ₂ Bloodgas
Partial Thromboplastin Time	Glucose
Prothrombin Time (Quick)	Troponin T
INR (Standardized Quick)	

Table A.2: Detailed information about hyperparameter search ranges. *: we fixed 10 Monte Carlo samples according to [Futoma et al. \(2017a\)](#). **: the MGP-RNN baseline was presented with a batch-size of 100, thus we did not enforce our range on this baseline ([Futoma et al., 2017a](#)).

Model	Hyperparameter	Lower bound	Upper bound	Sampling distribution
All models	learning rate	5e−4	5e−3	log uniform
	Monte Carlo Samples		10*	not applicable
MGP-TCN Raw-TCN	batch size	10	40	uniform
	temporal blocks	4	9	uniform
	filters per layer	15	90	uniform
	filter width	2	5	uniform
	dropout	0	0.1	uniform
	L_2 -regularization	0.01	100	log uniform
MGP-RNN	batch size		100**	not applicable
	layers	1	3	uniform
	hidden units per layer	20	150	uniform
	L_2 -regularization	1e−4	1e−3	log uniform

A.5. Hyperparameter Search

Differentiable models For the differentiable models *MGP-RNN*, *MGP-TCN*, and *Raw-TCN* an extensive hyperparameter search based on Bayesian optimization was performed using the `scikit-optimize` package ([scikit-optimize contributors, 2018](#)). More precisely, we relied on a Gaussian Process to model the AUPRC of the models dependent on the hyperparameters. The models were then trained at the hyperparameter values that matched one of the randomly-selected criteria *largest confidence bounds*, *largest expected improvement* and *highest probability of improvement* according to the Gaussian Process prior. A total of ten initial evaluations were performed at random according to the hyperparameter search space, followed by ten evaluations according to the Gaussian Process prior. During the hyperparameter search phase, the *MGP-RNN* model was trained for 5 epochs over the complete dataset—since we observed fast convergence behavior—while the TCN based model were trained 15 and 100 epochs for the *MGP-TCN* and the *Raw-TCN* model respectively.

DTW-KNN classifier The performance of the DTW-KNN classifier was evaluated on the same validation dataset as the other models for $k \in \{1, 3, \dots, 13, 15\}$, while relying on the training dataset with 0 hours before Sepsis onset. The k value yielding the best AUPRC on the validation dataset was selected for subsequent evaluation on the testing dataset. Similar to the other classifiers, we do not “refit” the k -nearest neighbors classifier by removing data from the training dataset to generate the horizon plots.

A.6. Additional Information on Scaling

Li and Marlin (2016) demonstrated that the MGP adapter framework is dominated by drawing from the MGP distribution. Thus, in our case, inverting $\Sigma_{\mathbf{i}} \in \mathbb{R}^{D \cdot T_i \times D \cdot T_i}$ has a computational complexity of $\mathcal{O}(D^3 \cdot T_i^3)$ (Golub and Van Loan, 2012). Notably, classifying a patient *only* depends on the length of the current patient time series and the number of tasks/variables; it does *not* depend on the number of patients in the training dataset; that is, it has a complexity of $\mathcal{O}(1)$ in the number of patients.

By contrast, while a naive implementation of DTW-KNN has a very low training complexity ($\mathcal{O}(1)$, due to its character as a “lazy learner”), the complexity at prediction time is very high. To classify a single instance, the k -nearest neighbors classifier requires the distances of the instance to all N training points. Moreover, the runtime of a single distance computation using dynamic time warping (DTW) is $\mathcal{O}(DT^2)$, where D represents the number of channels in the time series and T is the length of the shorter time series. Overall, this leads to a runtime complexity of $\mathcal{O}(NDT^2)$ for a *single* prediction step, which can quickly become infeasible for large-sized health record datasets, especially if online predictions are desired. Consequently, already for $N \geq D^2T$, the complexity of the DTW-KNN approach will exceed that of the MGP-TCN. Furthermore, the cubic complexity in the prediction step can be ameliorated by using faster approximation schemes (Li and Marlin, 2016).

A.7. Supplementary Results

To enforce a maximum time of two hours per call for each method (in order to make different hyperparameter searches on several folds feasible), MGP-RNN trains for 5 epochs, MGP-TCN for 15 epochs, and Raw-TCN for 100 epochs. We applied early stopping based on validation AUPRC with patience = 5 epochs.

Moreover, in an auxiliary setup, to guard against underfitting, we use the best parameter setting of each deep model and retrain each model for a prolonged period (50 epochs for both MGP-based models, 100 epochs for Raw-TCN) using early stopping based on validation AUPRC with patience = 10 epochs. As shown in Figure A.7, the deep models exhibit a similar tendency as in Figure 4, with the exception of the Raw-TCN showing high variability (due to one fold with favorable performance). Also, the MGP-based models exhibit a slight drop in performance. This may indicate that in terms of epochs, they converge earlier and overfit earlier.

A.8. Additional Information for the Horizon Analysis

When creating the horizon analysis, we observed that the current MGP-RNN implementation (using a Lanczos iteration) requires the minimal number of observed measurements of a patient to be at least the number of Monte Carlo samples (i.e. 10 in our case). Hence, we performed an on-the-fly masking of encounters that did *not* satisfy this criterion; for comparability, we applied it to all models. Table A.3 details the patient counts obtained after masking. Additionally, to be able fit all methods into memory, we removed a single outlier encounter consisting of more than 10K measurements. Notably, because we did not refit the models on each horizon (as this would answer a different question), with increasing prediction horizon, the slight decrease in sample size should not bias the mean performance but rather scale up the error bars.




Figure A.1: The auxiliary setup as computed for the full horizon. We retrained the best parameter settings of all deep model for a prolonged time to guard against underfitting. Due to a slightly decreased performance, the MGP-based models show a tendency to overfit when training for 50 epochs, whereas the Raw-TCN shows high variance between the random splits.

Table A.3: Patients count after applying masking which was required for making the MGP-RNN baseline work.

Horizon	Split Fold	Train			Validation			Test		
		0	1	2	0	1	2	0	1	2
0		4953	4950	4950	618	618	619	617	620	619
1		4951	4947	4947	618	618	619	616	620	619
2		4943	4941	4937	616	617	619	616	617	619
3		4933	4933	4927	614	615	617	615	614	618
4		4913	4917	4910	613	611	616	613	611	613
5		4832	4827	4830	602	605	602	598	600	600
6		4587	4580	4565	566	570	581	567	570	574
7		4073	4061	4056	503	508	510	497	504	507

A.9. Implementation Details & Runtimes

For maximal reproducibility, we embedded our method and all comparison partners in the `sacred v0.7.4` environment (Greff *et al.*, 2017). A local installation of the MIMIC-III database was done with PostgreSQL 9.3.22. The queries to extract the data from the database are based on queries in the public `mimic-code` repository (Johnson *et al.*, 2018). However, to extract the hourly-resolved sepsis label, we had to implement an entire query pipeline on top of the original code (Johnson *et al.*, 2018). For the MGP module, we included the code implemented by (Futoma *et al.*, 2017a) with minor changes. For

the TCNs, we extend the TensorFlow implementation of Lee (2018). We further apply *gradient checkpointing* (Chen *et al.*, 2016) for all neural network models in order to permit training in a typical GPU setup. The DTW-KNN classifier was implemented using the libraries `tslearn v0.1.26` (Tavenard, 2017) for dynamic time warping and `scikit-learn v0.20.2` (Pedregosa *et al.*, 2011) for the k -nearest neighbor classifier. We implemented both our proposed methods as well as all comparison partners in Python 3. All experiments were performed on a Ubuntu 14.04.5 LTS server with 2 CPUs (Intel[®] Xeon[®] E5-2620 v4 @ 2.10GHz), 8 GPUs (NVIDIA[®] GeForce[®] GTX 1080), and 128 GiB of RAM. However, for the deep learning models, we exclusively used single GPU processing. Supplementary Table A.4 depicts the runtimes of all methods.

Table A.4: Training runtimes (for the RNN/TCN methods, this includes the sum of all three splits, whereas for DTW-KNN distances were computed only once).

Method	Raw-TCN	MGP-RNN	MGP-TCN	DTW-KNN
Runtime	49.7 h	74.8 h	73.4 h	136.9 h

References

- Bonilla, E. V. *et al.* (2008). Multi-task gaussian process prediction. In *Advances in Neural Information Processing Systems*, pages 153–160.
- Chen, T. *et al.* (2016). Training deep nets with sublinear memory cost. *arXiv preprint arXiv:1604.06174*.
- Futoma, J. *et al.* (2017). Learning to detect sepsis with a multitask Gaussian process RNN classifier. In *International Conference on Machine Learning*, pages 1174–1182.
- Golub, G. H. and Van Loan, C. F. (2012). *Matrix computations*, volume 3. JHU Press.
- Greff, K. *et al.* (2017). The Sacred Infrastructure for Computational Research. In Katy Huff, David Lippa, Dillon Niederhut, and M. Pacer, editors, *Proceedings of the 16th Python in Science Conference*, pages 49–56.
- Johnson, A. E. *et al.* (2018). The MIMIC Code Repository: enabling reproducibility in critical care research. *Journal of the American Medical Informatics Association*, **25**(1), 32–39.
- Kingma, D. P. and Welling, M. (2014). Auto-encoding variational bayes. *Proceedings of the 2nd International Conference on Learning Representations (ICLR)*.
- Lee, C. (2018). Solving sequential MNIST with Temporal Convolutional Networks (TCNs). <https://colab.research.google.com/drive/11a331W7FQV1RicpfzyLq9H0SH1VSD4LE#scrollTo=b10XJFD-y-fS>.
- Li, S. C.-X. and Marlin, B. M. (2016). A scalable end-to-end Gaussian process adapter for irregularly sampled time series classification. In *Advances in Neural Information Processing Systems*, pages 1804–1812.

- Pedregosa, F. *et al.* (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, **12**, 2825–2830.
- scikit-optimize contributors, T. (2018). `scikit-optimize/scikit-optimize: v0.5.2`. <https://doi.org/10.5281/zenodo.1207017>.
- Seymour, C. W. *et al.* (2016). Assessment of clinical criteria for sepsis: For the Third International Consensus Definitions for Sepsis and Septic Shock (Sepsis-3). *Journal of the American Medical Association (JAMA)*, **315**(8), 762–774.
- Singer, M. *et al.* (2016). The Third International Consensus Definitions for Sepsis and Septic Shock (Sepsis-3). *Journal of the American Medical Association (JAMA)*, **315**(8), 801–810.
- Tavenard, R. (2017). `tslearn`: A machine learning toolkit dedicated to time-series data. <https://github.com/rtavenar/tslearn>.
- Vincent, J.-L. *et al.* (1996). The SOFA (sepsis-related organ failure assessment) score to describe organ dysfunction/failure. *Intensive Care Medicine*, **22**(7), 707–710.
- Williams, C. K. and Rasmussen, C. E. (2006). *Gaussian processes for machine learning*. MIT Press, Cambridge, MA, USA.