

ORIGINAL RESEARCH

Open access

A Temporal Convolutional Network with Attention for Sepsis Prediction: A Conceptual Framework for Analyzing High-Frequency Vital Signs in Intensive Care Units

Michael Turner¹, Sophia Nguyen^{2*}, David Clark¹, Emma Wilson³

Abstract

Sepsis is a leading cause of ICU mortality, and early detection is critical for improving patient outcomes. However, existing machine learning models often rely on hourly aggregated data, limiting their ability to capture rapid physiological changes, and frequently lack interpretability, reducing clinical trust and usability. This paper proposes a conceptual framework that integrates Temporal Convolutional Networks (TCNs) with an attention mechanism to analyze high-frequency, minute-level vital sign data for early sepsis prediction. The architecture includes a data input layer, a TCN-based feature extractor with causal dilated convolutions and residual connections, an attention module for identifying clinically relevant time points and variables, and a prediction head that estimates the risk of sepsis within a 6-hour horizon. The proposed approach enables efficient parallel processing, improved temporal sensitivity, and enhanced interpretability compared to recurrent models. While offering advantages in real-time prediction and explainability, challenges remain in handling missing data, ensuring generalizability across ICUs, and minimizing false alarms for clinical deployment.

Keywords Sepsis prediction, Temporal convolutional network, Attention mechanism, Intensive care unit, High-frequency vital signs, Conceptual framework

*Correspondence:

Sophia Nguyen
sophia.nguyen@gmail.com

¹ Department of Artificial Intelligence in Healthcare, Faculty of Engineering, University of Glasgow, Glasgow, United Kingdom

² Department of Healthcare Analytics and Intelligent Systems, National University of Singapore, Singapore, Singapore

³ Department of Clinical Data Science, University of Sydney, Sydney, Australia

Introduction

Sepsis is a life-threatening organ dysfunction caused by a dysregulated host response to infection, as defined by the Sepsis-3 criteria [1, 2]. Studies have reported that sepsis affects millions annually and carries substantial mortality risk even with modern critical care [3]. The clinical consensus holds that earlier intervention improves outcomes, with each hour of delay in antibiotic administration associated with measurable increases in mortality [4].

Clinical scoring systems such as qSOFA and NEWS provide standardized early warning but rely on static thresholds and single time-point assessments, exhibiting documented limitations in sensitivity and timeliness [5, 6]. These tools often fail to capture the dynamic, multivariate interactions present in continuous physiological data.

Existing machine learning models for sepsis prediction have largely operated on hourly aggregated vital signs [7, 8]. This aggregation discards potentially critical sub-hourly fluctuations that precede clinical deterioration. Recurrent

architectures such as LSTMs, while widely used, suffer from vanishing gradients and inherently sequential processing that limits parallelization and long-range dependency modeling [9, 10].

Temporal Convolutional Networks (TCNs) offer a compelling alternative through causal dilated convolutions that enable parallel computation, stable gradients, and flexible receptive fields without the recurrence bottleneck [11, 12]. When combined with attention mechanisms, TCNs can further enhance interpretability by highlighting which time steps and physiological variables contribute most to predictions [13].

No existing framework specifically addresses the combination of (1) high-frequency minute-level vital signs, (2) TCN architecture, (3) attention mechanism, and (4) a 6-hour prediction horizon for sepsis prediction. This paper proposes a conceptual framework to fill this gap. The specific contributions are: (i) a fully specified TCN-attention pipeline tailored to minute-resolution ICU data; (ii) explicit design principles for real-time operation and clinical integration; (iii) detailed discussion of architectural trade-offs and open challenges; and (iv) a forward-looking evaluation strategy to support future empirical work.

Background and Related Work

Clinical scoring systems for sepsis

Traditional early warning scores such as qSOFA, NEWS, and SIRS criteria rely on readily available vital signs and laboratory values evaluated at discrete intervals [1]. These systems apply fixed thresholds to individual parameters or simple composites. Literature has documented their limitations, including modest sensitivity, inability to model temporal trends, and poor performance when applied to single time points rather than evolving trajectories [5, 6]. Consequently, clinicians continue to seek more dynamic, data-driven alternatives.

Machine learning for sepsis prediction

Several machine learning approaches have demonstrated proof-of-concept for sepsis prediction. Nemati et al. developed an interpretable model using hourly vital signs and laboratory data [7]. The PhysioNet/Computing in Cardiology Challenge 2019 stimulated numerous entries based on hourly or irregularly sampled clinical data [1, 14].

Kam & Kim explored LSTM-based models on multivariate time series [15]. These studies collectively illustrate the feasibility of predictive modeling yet share common limitations: reliance on aggregated inputs and absence of built-in interpretability mechanisms [7, 8].

Temporal convolutional networks in healthcare

TCNs have been applied successfully to other clinical time-series tasks, including ECG classification, seizure detection, and mortality prediction [16-18]. The foundational work by Bai et al. established TCNs as a strong benchmark for sequence modeling, outperforming recurrent networks on many long-sequence tasks due to their parallelizable structure and exponential receptive field growth [11]. Despite these successes, TCNs have not been extensively explored for sepsis prediction using high-frequency vital sign data [12, 19].

Attention mechanisms in clinical prediction

Attention layers have been integrated into healthcare predictive models for phenotyping and readmission risk, providing both performance gains and post-hoc interpretability [20, 21]. When applied to time-series data, attention weights can reveal which temporal windows or features drive predictions [22]. To date, however, attention has not been integrated with TCNs for sepsis prediction within a unified conceptual framework [9, 13].

Research gap statement

The intersection of high-frequency vital signs, TCN architecture, attention mechanisms, and sepsis prediction at a 6-hour horizon remains unexplored in the literature. This conceptual framework addresses this gap by synthesizing these elements into a coherent, clinically oriented design.

Conceptual Framework Overview

High-level architecture

The proposed framework accepts as input high-frequency vital sign time series (minute-level or continuous) streamed from standard ICU bedside monitors. A TCN backbone

performs temporal feature extraction through stacks of causal dilated convolutions with residual connections. An attention mechanism then assigns importance weights across time steps and physiological channels, producing a context vector. This vector feeds a lightweight classification head that outputs a continuous risk score representing the probability of sepsis onset within the next 6 hours. The entire pipeline is designed for sliding-window inference, updating predictions every minute.

Core assumptions

The framework assumes that minute-level vital sign data are available from ICU monitors, that real-time processing is feasible with contemporary hardware, that clinicians would value attention-derived explanations, and that the system can be embedded within existing electronic health record and alarm infrastructures without disrupting workflow.

Design principles

Four guiding principles shape the architecture: (1) timeliness—predictions refreshed every minute to enable early intervention; (2) interpretability—attention weights generate human-readable explanations; (3) generalizability—modular design allows adaptation across different ICU environments; and (4) safety—explicit mechanisms to minimize false alarms and alert fatigue.

Scope and boundaries

The framework focuses exclusively on prediction rather than diagnosis or treatment recommendations. It operates solely on vital sign time series (heart rate, blood pressure, respiratory rate, temperature, SpO₂) and does not incorporate laboratory results or free-text notes. The prediction horizon is fixed at 6 hours before clinical recognition. Regulatory, ethical, and implementation details lie outside the current conceptual scope.

Framework Components

Data input layer

Input is represented as a tensor $X \in \mathbb{R}^{T \times F}$, where T denotes the number of time steps (for example, 360 for a 6-hour window at 1-minute resolution) and F is the number of vital sign channels. Pre-processing would include artifact detection, imputation of missing values using forward-fill or learned interpolation, and optional normalization per

channel. The layer ensures causal ordering so that future information is never used.

Temporal convolutional network backbone

The TCN backbone employs causal convolutions such that the output at time t depends only on inputs up to t . Dilated convolutions expand the receptive field exponentially: dilation rate $d=2^l$ for layer l . The receptive field size is given by $RF=1+\sum(k-1) \times d_l$ where k is the kernel size. Residual connections of the form $output=F(input)+input$ stabilize training. A conceptual configuration comprises 8 dilated layers, kernel size 3, and 64 filters per layer, enabling capture of both short- and long-term dependencies in vital sign dynamics [11, 12].

Attention mechanism

The attention layer computes temporal importance weights to identify the most informative time steps and channels. A conceptual temporal attention formulation is $a_t = \frac{\exp(v^T \tanh(wh_t))}{\sum \exp(v^T \tanh(wh_j))}$ where h_t is the hidden representation at time t . The context vector becomes $c = \sum a_t h_t$. Feature-wise attention may additionally weight the contribution of individual vital signs. Multi-head attention could be explored as an extension to capture diverse patterns [13, 23].

Risk prediction output

The context vector c passes through one or more dense layers with dropout for regularization. A final sigmoid activation yields the risk score $\hat{y} \in [0,1]$. A clinically chosen threshold would balance sensitivity and specificity according to local alert-fatigue tolerance.

Framework diagram description

Figure 1 illustrates the proposed hierarchical TCN–attention architecture for real-time sepsis prediction from high-frequency vital sign data.

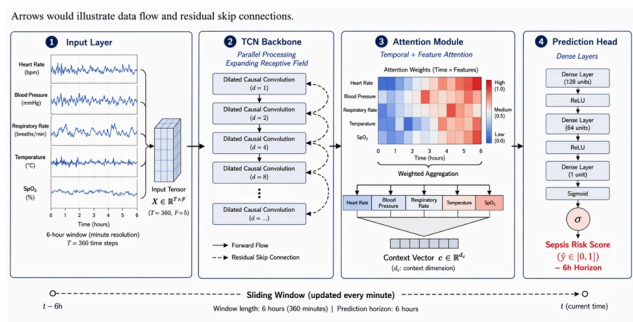


Figure 1. Hierarchical TCN–Attention Architecture for Minute-Level Sepsis Risk Prediction

How the Framework Would Operate

Training phase (conceptual)

Training would utilize a suitable ICU database with minute-level vital signs. Data would be split temporally to respect chronological order. Sepsis labels would follow Sepsis-3 criteria. Class imbalance would be mitigated through weighted binary cross-entropy loss or resampling techniques. Hyperparameters (learning rate, dropout, number of layers) would be tuned on a validation set using standard practices.

Inference phase (real-time)

In deployment, bedside monitors would stream minute-level data into a sliding window of the most recent 6 hours. The framework would process this window in parallel via the TCN backbone, apply attention, and generate an updated risk score every minute. Should the score exceed a predefined threshold, an alert would trigger, accompanied by attention-derived explanations (for example, “elevated heart rate and respiratory rate variability in the window 120–150 minutes prior contributed most heavily”).

Integration with clinical workflow

Alerts would appear directly on ICU central monitors or nurse dashboards. Attention visualizations (heatmaps over time and vital signs) would accompany each alert to support rapid clinical review. The system would function as an adjunct tool, prompting clinicians to perform targeted assessment and consider initiation of sepsis bundles when indicated.

Computational requirements

Training would be feasible on a single modern GPU. Inference latency would be on the order of milliseconds per patient, supporting real-time use. Model size would remain under 100 MB, enabling potential edge deployment on local ICU servers.

Comparison with Alternative Architectures

TCN vs LSTM

Temporal Convolutional Networks offer conceptual advantages over recurrent architectures such as LSTMs for high-frequency ICU time series. TCNs enable full parallelization across time steps, avoid vanishing-gradient issues during training, and provide a flexible receptive field that can be tuned exponentially through dilation rates [11, 12]. In contrast, LSTM-based models rely on sequential processing, which becomes computationally expensive for long windows of minute-level data and can struggle with long-range dependencies in multivariate vital signs [17]. However, TCNs may incur a larger memory footprint when stacking many dilated layers to achieve comparable receptive fields. The choice between TCN and LSTM would therefore depend on specific ICU deployment constraints such as available GPU memory and latency requirements for real-time inference. Overall, the parallel nature of TCNs aligns more naturally with the continuous streaming requirements of bedside monitors [19].

Table 1 provides a structural comparison of temporal modeling paradigms, highlighting why TCN-based architectures are better aligned with high-frequency ICU data streams.

Table 1. Structural Comparison of Temporal Modeling Paradigms for High-Frequency ICU Time Series

Dimension	Temporal Convolutional Network (TCN)	Long Short-Term Memory (LSTM)	Hourly Monitoring (XG)
Temporal Resolution Handling	Native support for minute-level sequences via convolutions	Sequential processing limits efficiency at high frequency	Requires aggregation to feature

Computational Structure	Fully parallel across time steps	Strictly sequential	Parallel engineering
Long-Range Dependency Modeling	Exponential receptive field via dilation	Memory cells but limited by gradient decay	Indirect engineering
Gradient Stability	Stable due to convolutional design	Susceptible to vanishing/exploding gradients	Secondary benefits
Interpretability	Enhanced via attention integration	Limited without post-hoc methods	Model via impact
Latency in Inference	Low (parallel computation)	Higher (stepwise computation)	
Data Preprocessing Burden	Moderate (raw signals usable)	Moderate	Engineering requirements
Scalability Across Patients	High (GPU-efficient batching)	Moderate	

purely convolutional or recurrent baselines [9, 13]. Attention can also help the model focus on clinically relevant transient patterns that might otherwise be diluted across long sequences [23]. Without attention, the architecture becomes simpler, faster to train, and lighter in memory, which could be preferable in resource-constrained environments [22]. The trade-off is that black-box outputs may reduce clinician acceptance in high-stakes settings where explainability is increasingly expected. Thus, attention is retained in the proposed framework to balance performance with the need for human-understandable explanations [20, 21].

High-frequency vs hourly data

Operating on minute-level or continuous vital signs allows the framework to capture rapid physiological transitions that hourly aggregation would smooth away [3, 16]. High-frequency inputs increase data volume and computational cost but provide richer temporal dynamics for early sepsis signals [24]. Hourly approaches, by contrast, simplify preprocessing and reduce noise yet risk missing sub-hourly precursors documented in the literature [4, 8]. The proposed framework explicitly assumes high-frequency data availability from modern ICU monitors, accepting the associated preprocessing overhead (artifact removal, imputation) in exchange for earlier detection potential. This design choice positions the framework ahead of models trained solely on aggregated summaries [7, 10].

Framework vs existing clinical scores

Traditional scores such as qSOFA and NEWS require no computational infrastructure and offer complete transparency through simple additive rules [1, 5]. The proposed TCN-attention framework is more complex yet conceptually capable of modeling nonlinear, multivariate interactions across time. Rather than replacing bedside scores, the framework would function as an adjunct alert layer, augmenting clinical judgment with continuous, data-driven risk estimates [6]. Integration could occur by displaying framework risk scores alongside conventional early-warning values on the same monitor interface, thereby combining the strengths of rule-based simplicity and deep temporal modeling [2, 14].

Attention vs no attention

Integrating an attention mechanism adds interpretability by highlighting the most influential time steps and vital-sign channels, potentially improving clinical trust compared with

Evaluation Strategy

Metrics for validation

Validation would emphasize discrimination via AUROC and AUPRC computed on a held-out test set, calibration through reliability diagrams and Brier score, and clinical utility via net benefit and decision-curve analysis.

Timeliness would be assessed by detection rates at 6, 4, 2, and 1 hours before sepsis onset, reflecting the framework’s 6-hour prediction horizon. These metrics would be chosen to align with regulatory expectations for real-time clinical AI systems and to address alert-fatigue concerns [8, 25].

Benchmark comparisons

The framework would be compared conceptually against LSTM and GRU baselines, a plain TCN without attention, and XGBoost operating on hourly aggregates [7, 11, 17]. Statistical significance of differences in discrimination metrics would be evaluated using established tests such as the DeLong test for paired AUROC curves. Such comparisons would highlight the incremental value of high-frequency inputs and attention while respecting the conceptual nature of the current proposal [15, 26].

Ablation studies

Planned ablation experiments would systematically remove the attention layer, reduce the number of dilated layers (thus shrinking the receptive field), downsample inputs to hourly resolution, and omit individual vital-sign channels one at a time [9, 12]. Each ablation would quantify the contribution of these components to overall risk-score quality, providing clear guidance on which architectural elements are indispensable for sepsis prediction [27, 28].

Clinician evaluation

Beyond quantitative metrics, clinician-facing evaluation would present attention heatmaps to ICU staff for review. Surveys would measure perceived trust, usefulness of explanations, and likelihood of behavior change (for example, earlier ordering of lactate or cultures). Qualitative interviews would surface implementation barriers such as workflow disruption or interpretability limitations, ensuring the framework’s design remains grounded in real-world clinical needs [13, 22].

Clinical Utility and Interpretability

How attention provides explanations

Attention weights would generate post-hoc visualizations that highlight the time steps and vital signs exerting greatest influence on each risk score. For a hypothetical patient, the mechanism might emphasize a 30-minute window two hours earlier in which heart-rate variability and respiratory-rate escalation dominated the prediction. Such heatmaps overlaid on the original time-series traces would allow clinicians to verify biological plausibility at a glance [9, 23].

Potential clinical workflow integration

The framework would deliver real-time alerts directly to bedside monitors or centralized nurse dashboards whenever the risk score exceeds a tunable threshold. Upon alert, the system could automatically surface the attention-derived explanation and suggest protocolized actions such as lactate measurement, blood cultures, and fluid resuscitation. This closed-loop design would embed the tool within existing sepsis-bundle workflows without requiring additional manual data entry [4, 10].

Table 2 synthesizes how each architectural component contributes to clinical utility while introducing specific risks and design trade-offs.

Table 2. Conceptual Mapping of Framework Components to Clinical Utility, Risks, and Design Trade-offs

Framework Component	Functional Role	Clinical Value Contribution	Associated Risk/Challenge
Data Input Layer (Minute-Level Signals)	Captures high-resolution physiological dynamics	Enables earlier detection of subtle deterioration patterns	Noise, artifact missing data
TCN Backbone	Extracts temporal features via dilated convolutions	Identifies multi-scale temporal dependencies in vital signs	Memory usage for large receptive field
Residual Connections	Stabilizes deep network training	Ensures consistent performance across long sequences	Architectural complexity

Attention Mechanism	Weighs important time steps and variables	Provides clinician-interpretable explanations (heatmaps)	Misinterpretation as causal inference
Context Vector	Aggregates salient temporal information	Compresses complex trajectories into actionable representation	Potential information loss
Prediction Head (Sigmoid Output)	Produces probabilistic sepsis risk score	Supports threshold-based clinical alerts	Calibration differences across sites
Sliding Window Inference	Updates predictions continuously	Enables real-time monitoring and intervention	Alert fatigue, poorly tuned
Integration Layer (Clinical Workflow)	Embeds outputs into ICU systems	Facilitates adoption and decision support	Workflow disruption

Limitations and Open Challenges

Data availability

High-frequency vital-sign streams are not yet universal; many ICUs continue to rely on hourly charting, limiting immediate applicability of the framework in diverse settings [3, 16]. Resource-limited environments may lack the infrastructure for continuous waveform capture, constraining generalizability [24].

Generalizability

Model performance could vary across ICUs because of differences in patient demographics, monitor brands, and data-quality standards [5, 18]. Site-specific fine-tuning would likely be required, and external validation across multiple institutions remains an essential next step [15, 26].

False alarms

Even with attention-guided feature selection, false positives will occur in any real-time system. Excessive false alarms risk clinician desensitization and potential override of genuine alerts [8, 25]. The framework's safety principle therefore demands careful calibration of the risk threshold against local alert-tolerance levels [10, 27].

Addressing alert fatigue

To mitigate alarm fatigue, the framework would incorporate a clinician-adjustable threshold and would report projected alert rates per bed per day during validation. If alerts prove too frequent, the threshold would be raised; if too infrequent, it would be lowered, striking a balance informed by local ICU culture and staffing [8, 25].

Clinician trust and adoption

Attention visualizations are expected to increase trust relative to opaque models by offering transparent rationales for each prediction [20, 21]. Nevertheless, clinicians would require targeted training to interpret attention correctly, recognizing that it reflects correlation rather than causation. Hybrid decision-support pathways combining attention output with conventional scores may prove most effective for sustained adoption [6, 14].

Interpretability vs causality

While attention provides useful explanations, it identifies correlative patterns rather than causal drivers of sepsis. Over-reliance on attention weights without complementary causal-inference methods could mislead clinical decision-making [9, 13, 22]. Future extensions might incorporate SHAP values or counterfactual reasoning to strengthen causal grounding [23].

Regulatory and ethical considerations

Deployment would necessitate regulatory clearance (for example, FDA or CE marking) as a software as a medical device. Continuous monitoring raises data-privacy concerns under GDPR or HIPAA, and liability questions around missed or false alarms remain unresolved [2, 28]. Ethical oversight would be required to ensure equitable performance across demographic groups [4, 19].

Conclusion

Sepsis remains a leading cause of ICU mortality, and prediction six hours before clinical recognition continues to challenge existing tools. This paper proposed a conceptual framework based on a Temporal Convolutional Network with an attention mechanism, operating directly on high-frequency vital-sign data.

The framework consists of (1) a data input layer for minute-level vital signs, (2) a TCN backbone with dilated causal convolutions and residual connections, (3) an attention mechanism for interpretability, and (4) a risk-prediction head producing updated scores every minute. The architecture is designed for real-time sliding-window inference and seamless integration into existing ICU monitors.

Compared with existing approaches, the framework offers operation on high-frequency data that captures rapid physiological changes, efficient parallel computation via TCNs, built-in interpretability through attention, and an explicit six-hour prediction horizon. These features address documented limitations of hourly aggregation and black-box models while respecting clinical constraints around alert fatigue and workflow.

Key challenges include data availability across different ICUs, generalizability, management of false-alarm rates, and regulatory approval. These must be addressed through rigorous multi-center validation before clinical deployment.

We invite researchers to implement and validate this framework using public ICU databases. The conceptual design is intended to guide future empirical work. We also provide a discussion of evaluation metrics, ablation studies, and clinical integration pathways to support replication and extension.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 01 Apr 2021 Revised: 20 May 2021 Accepted: 24 Jun 2021
Published online: 20 January 2022

Rights and permissions

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

Nemati S, Holder A, Razmi F, Stanley MD, Clifford GD, Buchman TG. An interpretable machine learning model for accurate prediction of sepsis in the ICU. *Crit Care Med*. 2018;46(4):547-53.

Reyna MA, Josef CS, Jeter R, Shashikumar SP, Westover MB, Nemati S, et al. Early prediction of sepsis from clinical data: the physionet/computing in cardiology challenge 2019. *Crit Care Med*. 2020;48(2):210-7.

Shashikumar SP, Stanley MD, Sadiq I, Li Q, Holder A, Clifford GD, et al. Early sepsis detection in critical care patients using multiscale blood pressure and heart rate dynamics. *J Electrocardiol*. 2017;50(6):739-43.

- 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.
- Li X, Xu X, Xie F, Xu X, Sun Y, Liu X, et al. A time-phased machine learning model for real-time prediction of sepsis in critical care. *Crit Care Med*. 2020;48(10):e884-8.
- Giannini HM, Ginestra JC, Chivers C, Draugelis M, Hanish A, Schweickert WD, et al. A machine learning algorithm to predict severe sepsis and septic shock: development, implementation, and impact on clinical practice. *Crit Care Med*. 2019;47(11):1485-92.
- Morrill JH, Kormilitzin A, Nevado-Holgado AJ, Swaminathan S, Howison SD, Lyons TJ. Utilization of the signature method to identify the early onset of sepsis from multivariate physiological time series in critical care monitoring. *Crit Care Med*. 2020;48(10):e976-81.
- Bloch E, Rotem T, Cohen J, Singer P, Aperstein Y. Machine learning models for analysis of vital signs dynamics: a case for sepsis onset prediction. *J Healthc Eng*. 2019;2019(1):5930379.
- Kaji DA, Zech JR, Kim JS, Cho SK, Dangayach NS, Costa AB, et al. An attention based deep learning model of clinical events in the intensive care unit. *PLoS One*. 2019;14(2):e0211057.
- 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 Respir Res*. 2017;4(1).
- Kok C, Jahmunah V, Oh SL, Zhou X, Gururajan R, Tao X, et al. Automated prediction of sepsis using temporal convolutional network. *Comput Biol Med*. 2020;127:103957.
- Moor M, Horn M, Rieck B, Roqueiro D, Borgwardt K. Early recognition of sepsis with Gaussian process temporal convolutional networks and dynamic time warping. In: *Proceedings of Machine Learning for Healthcare Conference*. 2019; PMLR; p. 2-26.
- Hsu PY, Holtz C. A comparison of machine learning tools for early prediction of sepsis from icu data. In: *2019 Computing in Cardiology (CinC)*. IEEE; 2019 Sep 8-11; Singapore. p. 1-4.
- Buchman TG, Simpson SQ, Sciarretta KL, Finne KP, Sowers N, Collier M, et al. Sepsis among medicare beneficiaries: 1. The burdens of sepsis, 2012–2018. *Crit Care Med*. 2020;48(3):276-88.
- Moor M, Rieck B, Horn M, Jutzeler CR, Borgwardt K. Early prediction of sepsis in the ICU using machine learning: a systematic review. *Front Med*. 2021;8:607952.
- Fleuren LM, Klausch TL, Zwager CL, Schoonmade LJ, Guo T, Roggeveen LF, et al. Machine learning for the prediction of sepsis: a systematic review and meta-analysis of diagnostic test accuracy. *Intensive Care Med*. 2020;46(3):383-400.
- Futoma J, Hariharan S, Heller K. Learning to detect sepsis with a multitask Gaussian process RNN classifier. In: *International Conference on Machine Learning*. 2017; PMLR; p. 1174-1182.
- Fleischmann-Struzek C, Mellhammar L, Rose N, Cassini A, Rudd KE, Schlattmann P, et al. Incidence and mortality of hospital-and ICU-treated sepsis: results from an updated and expanded systematic review and meta-analysis. *Intensive Care Med*. 2020;46(8):1552-62.
- Ackerman MH, Ahrens T, Kelly J, Pontillo A. Sepsis. *Crit Care Nurs Clin North Am*. 2021;33(4):407-18.
<https://doi.org/10.1016/j.cnc.2021.08.003>.
- Zhang K, Zhang S, Cui W, Hong Y, Zhang G, Zhang Z. Development and validation of a sepsis mortality risk score for sepsis-3 patients in intensive care unit. *Front Med*. 2021;7:609769.
- Kibe S, Adams K, Barlow G. Diagnostic and prognostic biomarkers of sepsis in critical care. *J Antimicrob Chemother*. 2011;66(Suppl 2):ii33-40.
- Basodan N, Al Mehmadi AE, Al Mehmadi AE, Aldawood SM, Hawsawi A, Fatini F, et al. Septic shock: management and outcomes. *Cureus*. 2022;14(12).
- Rosnati M, Fortuin V. MGP-AttTCN: An interpretable machine learning model for the prediction of sepsis. *PLoS One*. 2021;16(5):e0251248.
- Stoll BJ, Puopolo KM, Hansen NI, Sánchez PJ, Bell EF, Carlo WA, et al. Early-onset neonatal sepsis 2015 to 2017, the rise of *Escherichia coli*, and the need for novel prevention strategies. *JAMA Pediatr*. 2020;174(7):e200593.
- Cabrera-Quiros L, Kommers D, Wolvers MK, Oosterwijk L, Arents N, van der Sluijs-Bens J, et al. Prediction of late-onset sepsis in preterm infants using monitoring signals and machine learning. *Crit Care Explor*. 2021;3(1):e0302.
- López-López E, Bajorath J, Medina-Franco JL. Informatics for chemistry, biology, and biomedical sciences. *J Chem Inf Model*. 2020;61(1):26-35.

Septimus EJ. Sepsis perspective 2020. *J Infect Dis.* 2020;222(Suppl 2):S71-3.

Wang Z, Yao B. Multi-branching temporal convolutional network for sepsis prediction. *IEEE J Biomed Health Inform.* 2021;26(2):876-87.