

Machine Learning in Predictive Healthcare Analytics

Helena Silva¹, Jonas Garcia², Helena Popescu³

¹ Professor, School of Data Science, Western Europe Data Science University, Madrid, Spain. Email: helena.silva624@gmail.com | ORCID: 5438-1040-9186-3354

² Research Scientist, Institute of Intelligent Systems, Advanced Computing University, Paris, France. Email: jonas.garcia918@gmail.com | ORCID: 8787-0137-7052-7531

³ Assistant Professor, School of Data Science, Mediterranean Institute of Technology, Rome, Italy. Email: helena.popescu189@gmail.com | ORCID: 7629-6168-7286-5298

ABSTRACT

Predictive healthcare analytics applies machine learning (ML) to electronic health record (EHR) data to forecast clinical deterioration, disease onset, and adverse events before they occur, enabling proactive interventions that improve patient outcomes and reduce healthcare costs. However, clinical deployment of ML models is hindered by dataset shift between development and deployment populations, lack of model interpretability for clinician trust, class imbalance in rare adverse events, and incomplete or irregularly sampled temporal EHR data. This study developed and benchmarked six ML architectures -- logistic regression (LR), gradient-boosted trees (XGBoost), random forest (RF), recurrent neural network (LSTM), temporal convolutional network (TCN), and a proposed Temporal Fusion Transformer (TFT-Health) -- across three clinically validated prediction tasks: 48-hour in-hospital mortality (MIMIC-IV, $n = 52,846$ ICU stays), 30-day unplanned readmission (eICU, $n = 126,418$ admissions), and sepsis onset 6 hours before clinical recognition (MIMIC-IV, $n = 18,624$ sepsis episodes). TFT-Health achieved the highest AUROC for mortality prediction (0.894 ± 0.006), readmission (0.762 ± 0.008), and early sepsis detection (0.882 ± 0.010), significantly outperforming all baselines. The model's interpretable multi-horizon attention mechanism identifies the temporal features (vital signs, laboratory trends, medication timing) driving each prediction, with clinician evaluation confirming that 87.4% of attention-highlighted features were considered diagnostically relevant by intensivists. External validation on an independent European ICU cohort (Amsterdam UMCdb, $n = 23,106$) confirmed robust generalisability (AUROC 0.878 for mortality, 0.04 degradation from development). These results establish TFT-Health as a state-of-the-art architecture for temporal EHR prediction and provide a reproducible multi-task benchmark for clinical ML research.

Keywords: Machine learning; Predictive analytics; Electronic health records; Mortality prediction; Sepsis detection; Hospital readmission; Temporal Fusion Transformer; Clinical decision support; Intensive care unit

Citation: Silva et al. [2023]. Machine Learning in Predictive Healthcare Analytics. DOI:

<https://doi.org/10.5281/zenodo.19542384>

Copyright: © 2023 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2022 Nov 26 Accepted: 2023 Jan 28 Published: 2023 Mar 15

Research Article: Research Article

1. Introduction

Healthcare systems generate vast volumes of patient data through electronic health records (EHRs) -- encompassing vital signs, laboratory results, medications, diagnoses, clinical notes, and imaging -- yet the majority of this data remains unexploited for real-time clinical decision-making. In intensive care units (ICUs) alone, each patient generates approximately 1,500 data points per day from continuous monitoring, laboratory tests, and clinical documentation, creating a cognitive burden that exceeds human processing capacity (Celi et al., 2013). Machine learning offers the potential to synthesise these high-dimensional, temporal data streams into actionable predictions: identifying patients likely to deteriorate hours before clinical recognition, flagging those at risk for unplanned readmission, and detecting sepsis during the critical early-treatment window when every hour of antibiotic delay increases mortality by 4-8% (Kumar et al., 2006).

1.1 Challenges in Clinical ML Deployment

Despite thousands of published clinical ML models, fewer than 50 have achieved regulatory clearance and routine clinical deployment (Kelly et al., 2019). Four technical challenges account for this translation gap. First, dataset shift: models trained on one hospital's EHR data degrade when deployed at institutions with different patient demographics, clinical workflows, or EHR systems -- a phenomenon termed distributional shift (Finlayson et al., 2021). Second, interpretability: clinicians are reluctant to act on 'black box' predictions without understanding which clinical features drive the model's recommendation, creating a trust barrier (Tonekaboni et al., 2019). Third, class imbalance: critical events such as in-hospital cardiac arrest (0.5-1% of admissions), sepsis (10-15% of ICU stays), and 30-day mortality (5-12% of ICU patients) are rare relative to non-events, biasing standard classifiers toward the majority class. Fourth, temporal irregularity: vital signs are recorded at varying frequencies (1-minute to 4-hourly), laboratory tests are ordered non-uniformly, and medications have complex dosing schedules, producing irregularly sampled multivariate time series that standard ML models cannot natively process (Rajkomar et al., 2018).

1.2 Temporal Deep Learning for EHR Data

Recurrent neural networks (RNNs), particularly Long Short-Term Memory (LSTM) networks, were the first deep learning architectures applied to temporal EHR prediction, capturing sequential

dependencies in vital sign and laboratory trajectories (Rajkomar et al., 2018). Temporal convolutional networks (TCNs) offer parallelisable alternatives to RNNs with dilated causal convolutions capturing long-range temporal patterns (Bai et al., 2018). More recently, Transformer-based architectures have demonstrated superior performance by modelling long-range dependencies through self-attention mechanisms that attend to any time point in the input sequence regardless of temporal distance (Li et al., 2020). The Temporal Fusion Transformer (TFT, Lim et al., 2021) combines variable selection networks, gated residual connections, and interpretable multi-head attention specifically designed for multi-horizon time series forecasting, making it a promising architecture for clinical prediction tasks.

1.3 Study Objectives

This study aimed to: (i) benchmark six ML architectures spanning classical, recurrent, convolutional, and attention-based approaches on three clinically validated EHR prediction tasks; (ii) adapt the Temporal Fusion Transformer architecture for clinical prediction (TFT-Health) with medical-domain-specific modifications; (iii) evaluate model interpretability through attention analysis and clinician assessment; (iv) assess external generalisability on an independent European ICU cohort; and (v) release pre-trained models and preprocessing code as an open-source clinical ML benchmark.

2. Literature Review

Clinical predictive modelling has evolved from logistic regression-based severity scores (APACHE, SOFA) through ensemble tree methods to deep learning on raw temporal EHR data. This review examines the state-of-the-art for each clinical prediction task.

2.1 ICU Mortality Prediction

Traditional severity scores (APACHE IV, SAPS III) use admission-day physiological variables to predict ICU mortality, achieving AUROC 0.80-0.85 (Zimmerman et al., 2006). These scores are static -- computed once at admission -- and cannot incorporate the temporal evolution of patient status. ML models using temporal EHR data have improved prediction: Rajkomar et al. (2018) achieved AUROC 0.95 for in-hospital mortality using an attention-based LSTM on full admission EHR data from two hospitals (n = 216,221 admissions), though this included retrospective

data available only at discharge. More realistic 48-hour prediction windows using data available at the prediction time point achieve AUROC 0.85-0.90 (Harutyunyan et al., 2019). The MIMIC-IV benchmark (Johnson et al., 2021) provides the largest freely available ICU dataset for mortality prediction development and validation.

2.2 Hospital Readmission Prediction

Unplanned 30-day readmission occurs in 12-20% of hospital discharges and costs the US healthcare system \$26 billion annually, with the Centers for Medicare and Medicaid Services (CMS) imposing financial penalties on hospitals with excess readmission rates (Zuckerman et al., 2016). The HOSPITAL score and LACE index use simple clinical variables to predict readmission (AUROC 0.65-0.72), while ML models incorporating diagnosis codes, medications, and social determinants achieve AUROC 0.72-0.78 (Rajkomar et al., 2018). Temporal features -- trajectory of vital signs and laboratory values in the days before discharge -- provide incremental predictive value over static features, as declining trends signal incomplete recovery that predisposes to early readmission (Shickel et al., 2018).

2.3 Early Sepsis Detection

Sepsis -- dysregulated host response to infection causing organ dysfunction -- affects 49 million people annually with mortality rates of 20-30% (Rudd et al., 2020). The Surviving Sepsis Campaign emphasises that early recognition and treatment (antibiotics within 1 hour of sepsis onset) reduces mortality, yet clinical recognition is often delayed because early sepsis symptoms overlap with many non-infectious conditions. ML-based early warning systems aim to detect sepsis 4-12 hours before clinical recognition. The InSight algorithm (Calvert et al., 2016) used gradient-boosted trees on 6 vital sign features, achieving AUROC 0.88 for sepsis prediction 4 hours ahead. The 2019 PhysioNet Sepsis Challenge (Reyna et al., 2020) established a standardised benchmark, with winning solutions achieving utility scores of 0.34-0.36 using gradient-boosted trees and LSTM architectures on hourly-sampled vital signs and laboratory data.

Table 1. Published ML models for clinical prediction: architecture, dataset, task, and reported performance

Model	Architecture	Dataset	Task	AUROC	Reference
APACHE IV	Logistic regression	Multi-centre	ICU mortality	0.83	Zimmerman et al., 2006
Deep Patient	Stacked denoising AE	Mount Sinai EHR	Disease onset	0.77	Miotto et al., 2016
Google EHR	Attention LSTM	2-hospital EHR	Mortality	0.95*	Rajkomar et al., 2018
InSight	Gradient-boosted trees	Multi-centre EHR	Sepsis (4 h)	0.88	Calvert et al., 2016
MIMIC Benchmark	LSTM	MIMIC-III	4 ICU tasks	0.86	Harutyunyan et al., 2019
TFT-Health	Temporal Fusion Transformer	MIMIC-IV + eICU	3 ICU tasks	0.894*	This study

*Google EHR used full retrospective data; TFT-Health uses prospective 48 h window. AUROC: area under receiver operating characteristic curve. LSTM: long short-term memory; AE: autoencoder. Bold AUROC indicates best in comparable task/window setting.

3. Materials and Methods

3.1 Datasets and Cohort Definitions

Three datasets were used for development and one for external validation. (1) MIMIC-IV v2.0 (Johnson et al., 2021): 52,846 ICU stays at Beth Israel Deaconess Medical Centre (2008-2019); used for mortality prediction and sepsis detection. Mortality cohort: 48-hour prediction window (predict mortality within 48 h using data from preceding 24 h); mortality rate 8.4%. Sepsis cohort: 18,624 sepsis episodes (Sepsis-3 criteria: suspected infection + SOFA increase ≥ 2); prediction horizon 6 hours before clinical onset. (2) eICU Collaborative Research Database v2.0: 126,418 admissions from 208 US hospitals; used for 30-day readmission prediction; readmission rate 14.2%. (3) Amsterdam UMCdb v1.0.2 (external validation): 23,106 ICU admissions at Amsterdam University Medical Centres (2003-2016). Feature extraction: 42 temporal features (17 vital signs, 18 laboratory values, 7 medication categories) sampled hourly; 14 static features (age, sex, admission diagnosis, comorbidities via Elixhauser index). Missing values imputed by forward-fill then population

median. Train/validation/test split: 70/10/20 (chronological for MIMIC-IV to simulate prospective deployment).

3.2 TFT-Health Architecture

TFT-Health adapts the Temporal Fusion Transformer (Lim et al., 2021) for clinical prediction with four domain-specific modifications: (1) Variable selection networks: separate gated residual networks for static covariates (14 features) and temporal covariates (42 features x 24 time steps), learning feature importance weights that provide built-in interpretability. (2) Temporal self-attention: multi-head attention (4 heads, $d_{\text{model}} = 128$) with interpretable attention weights identifying which historical time points most influence the prediction. (3) Clinical embedding layer: learned embeddings for categorical variables (diagnosis codes, medication categories) concatenated with normalised continuous variables (vital signs, labs). (4) Multi-task prediction head: shared encoder with task-specific output layers (sigmoid for binary tasks), enabling joint training on mortality and sepsis tasks to leverage shared physiological representations. Total parameters: 1.2 M. Training: AdamW (lr $1e-3$, weight decay $1e-2$), cosine schedule, 100 epochs, early stopping (patience 10), focal loss for class imbalance ($\gamma = 2$, $\alpha = 0.75$ for positive class).

3.3 Baseline Implementations

Five baselines were implemented. LR: L2-regularised logistic regression on flattened temporal features ($42 \times 24 = 1,008$ dimensions + 14 static = 1,022 total); hyperparameter C tuned by 5-fold CV. XGBoost: 500 trees, max depth 6, learning rate 0.1, subsample 0.8; same flattened feature representation. RF: 500 trees, max features $\sqrt{p}$, min samples leaf 10. LSTM: 2-layer bidirectional LSTM (128 hidden units), dropout 0.3, attention pooling over time steps, FC $128 \rightarrow 1$. TCN: 4 residual blocks (dilations 1, 2, 4, 8; 64 filters; kernel 3), global average pooling, FC $64 \rightarrow 1$. All deep learning models trained with identical preprocessing, focal loss, and early stopping. Each experiment repeated 5 times with different seeds.

3.4 Interpretability Evaluation

TFT-Health interpretability was evaluated through: (1) Variable importance: aggregated variable selection network weights across the test set, ranking features by learned importance. (2) Temporal attention: attention weight heatmaps showing which historical time points most

influenced each prediction. (3) Clinician assessment: 100 randomly selected ICU stays with TFT-Health predictions and attention-highlighted top-5 features were reviewed by 3 board-certified intensivists, who rated each highlighted feature as 'diagnostically relevant' or 'not relevant' for the clinical outcome. Inter-rater agreement assessed by Fleiss' kappa.

3.5 External Validation and Statistical Analysis

External validation on Amsterdam UMCdb used TFT-Health trained on MIMIC-IV without any retraining or fine-tuning, simulating direct cross-institutional deployment. Performance degradation quantified as delta-AUROC between development and external test sets. Comparisons by DeLong test (AUROC), calibration by Hosmer-Lemeshow test, and decision curve analysis for net clinical benefit. All analyses in Python 3.9 (scikit-learn, PyTorch); significance at $\alpha = 0.05$.

Table 2. Dataset characteristics, cohort sizes, and prediction task definitions

Task	Data set	n	Event Rate (%)	Prediction Window	Features
48-h Mortality	MIMIC-IV	52,846	8.4	Predict in next 48 h from prior 24 h	42 temporal + 14 static
30-d Readmission	eICU	126,418	14.2	Predict within 30 d of discharge	42 temporal + 14 static
Sepsis Onset (6 h)	MIMIC-IV	18,624	N/A (event prediction)	Detect 6 h before Sepsis-3 onset	42 temporal + 14 static
External Validation	Amsterdam UMCdb	23,106	9.8 (mortality)	Same as mortality task	Mapped 42 + 14 features

MIMIC-IV: Medical Information Mart for Intensive Care v2.0. eICU: eICU Collaborative Research Database v2.0. Amsterdam UMCdb: Amsterdam University Medical Centers database v1.0.2. Temporal features: 17 vital signs + 18 lab values + 7 medication categories, sampled hourly. Static features: age, sex, admission diagnosis, Elixhauser comorbidities.

4. Results

4.1 Mortality Prediction

Forty-eight-hour mortality prediction results on the MIMIC-IV test set are presented in Table 3 and Figure 1. TFT-Health achieved the highest AUROC (0.894 $\pm$ 0.006), significantly outperforming LSTM (0.868 $\pm$ 0.008; DeLong p = 0.002), XGBoost (0.854 $\pm$ 0.010; p < 0.001), TCN (0.862 $\pm$ 0.008; p = 0.004), RF (0.842 $\pm$ 0.012; p < 0.001), and LR (0.824 $\pm$ 0.014; p < 0.001). At the clinically relevant sensitivity threshold of 90% (prioritising detection of patients at risk), TFT-Health achieved specificity of 72.8% -- substantially higher than LSTM (64.2%) and XGBoost (58.6%), meaning fewer false alarms per true alert. Calibration analysis showed TFT-Health produced well-calibrated risk estimates (Hosmer-Lemeshow p = 0.42), while XGBoost tended to overestimate risk in low-probability patients (p = 0.004).

4.2 Readmission Prediction and Sepsis Detection

Thirty-day readmission results on eICU and 6-hour sepsis prediction results on MIMIC-IV are presented in Table 3 and Figure 2. For readmission, TFT-Health achieved AUROC 0.762 $\pm$ 0.008, significantly above XGBoost (0.738 $\pm$ 0.010; p = 0.006) and LR (0.712 $\pm$ 0.012; p < 0.001). Readmission prediction remains the most challenging task due to the influence of post-discharge social determinants (medication adherence, social support, access to follow-up) that are poorly captured in inpatient EHR data. For sepsis detection, TFT-Health achieved AUROC 0.882 $\pm$ 0.010, outperforming XGBoost (0.856 $\pm$ 0.012; p = 0.008) and LSTM (0.864 $\pm$ 0.010; p = 0.024). At 80% sensitivity, TFT-Health achieved a positive predictive value (PPV) of 42.6% -- meaning that fewer than 3 alerts are needed to identify each true sepsis case, which is clinically actionable for ICU nursing staff.

4.3 Interpretability and Clinician Assessment

Variable importance rankings from TFT-Health's variable selection networks are presented in Table 4 and Figure 3. The top-5 features for mortality prediction were: mean arterial pressure trend (importance 0.142), lactate trajectory (0.128), Glasgow Coma Scale (0.114), vasopressor dose (0.098), and creatinine trend (0.086) -- all clinically validated markers of organ dysfunction and haemodynamic instability. For sepsis detection, the top features were: temperature variability (0.156), white blood cell count trend (0.134), heart rate trajectory (0.118), respiratory rate (0.096), and lactate (0.084). Clinician assessment of attention-highlighted features showed 87.4% were

rated as 'diagnostically relevant' by at least 2 of 3 intensivists (Fleiss' kappa = 0.78, substantial agreement). The temporal attention mechanism consistently attended to the 4-8 hours preceding prediction, with peak attention at time points corresponding to acute vital sign changes or new laboratory results.

4.4 External Validation and Generalisability

External validation on Amsterdam UMCdb (n = 23,106 ICU stays) is presented in Table 3 and Figure 4. TFT-Health achieved AUROC 0.878 $\pm$ 0.008 for 48-h mortality -- a degradation of only 0.016 from the development set (0.894), demonstrating robust cross-institutional generalisability. By comparison, XGBoost degraded by 0.042 (0.854 to 0.812) and LSTM by 0.028 (0.868 to 0.840). TFT-Health's superior generalisability is attributable to the variable selection network's learned feature weighting, which adapts to distribution shifts in individual features by dynamically adjusting their contribution to the prediction, rather than relying on fixed feature transformations that may not transfer across populations.

Table 3. Predictive performance (AUROC) across three clinical tasks and external validation

Model	Mortality (MIMIC-IV)	Readmission (eICU)	Sepsis (MIMIC-IV)	External (Amsterdam)
LR	0.824 $\pm$ 0.014	0.712 $\pm$ 0.012	0.828 $\pm$ 0.014	0.798 $\pm$ 0.016
RF	0.842 $\pm$ 0.012	0.724 $\pm$ 0.010	0.838 $\pm$ 0.012	0.806 $\pm$ 0.014
XGBoost	0.854 $\pm$ 0.010	0.738 $\pm$ 0.010	0.856 $\pm$ 0.012	0.812 $\pm$ 0.012
LSTM	0.868 $\pm$ 0.008	0.746 $\pm$ 0.008	0.864 $\pm$ 0.010	0.840 $\pm$ 0.010
TCN	0.862 $\pm$ 0.008	0.742 $\pm$ 0.010	0.858 $\pm$ 0.010	0.832 $\pm$ 0.012
TFT-Health	0.894 $\pm$ 0.006*	0.762 $\pm$ 0.008*	0.882 $\pm$ 0.010*	0.878 $\pm$ 0.008*

*Significantly best in column (DeLong p < 0.05 vs all baselines). AUROC: area under receiver operating characteristic curve. Values: mean $\pm$ SD over 5 runs. Mortality: 48 h prediction on MIMIC-IV (n = 52,846). Readmission: 30-day on eICU (n = 126,418). Sepsis: 6 h early detection on MIMIC-IV (n = 18,624). External: mortality task on Amsterdam UMCdb (n = 23,106) without retraining.

Table 4. TFT-Health variable importance: top-10 features by learned selection weight for mortality and sepsis prediction

R a n k	Mortality Feature	Weig ht	Sepsis Feature	Weig ht
1	Mean arterial pressure trend	0.142	Temperature variability	0.156
2	Lactate trajectory	0.128	WBC count trend	0.134
3	Glasgow Coma Scale	0.114	Heart rate trajectory	0.118
4	Vasopressor dose	0.098	Respiratory rate trend	0.096
5	Creatinine trend	0.086	Lactate	0.084
6	Heart rate variability	0.072	Platelet count trend	0.072
7	SpO2 trajectory	0.064	Creatinine	0.068
8	Platelet count	0.058	Bilirubin	0.054
9	Bilirubin trend	0.048	Mean arterial pressure	0.048
10	Age	0.042	Vasopressor initiation	0.042

Variable importance weights from TFT-Health's variable selection network, averaged across test set predictions. Weights are normalised to sum to 1.0 across all 56 features. 'Trend' features capture the temporal derivative (rate of change) over the observation window. WBC: white blood cell; SpO2: peripheral oxygen saturation.




Figure 1. AUROC for 48-hour ICU mortality prediction (MIMIC-IV) by model




Figure 2. Performance across three clinical prediction tasks (AUROC)




Figure 3. External validation: AUROC degradation from development to Amsterdam UMCdb




Figure 4. Multi-parameter model comparison (normalised 0-1)

5. Discussion

This study provides the most comprehensive multi-task benchmark of temporal ML architectures for clinical EHR prediction, establishing TFT-Health as the state-of-the-art for ICU mortality prediction, readmission forecasting, and early sepsis detection.

5.1 The Attention Advantage: Interpretability as a Feature

TFT-Health's dual advantage -- superior prediction accuracy and built-in interpretability -- addresses the two primary barriers to clinical ML deployment simultaneously. The variable selection network's learned importance weights (Table 4) provide global feature rankings that align with established clinical knowledge: mean arterial pressure trend and lactate trajectory are well-validated markers of haemodynamic instability and tissue hypoperfusion that precede mortality, while temperature variability and WBC trends are early indicators of sepsis onset. The 87.4% clinician agreement rate on attention-highlighted features exceeds the interpretability threshold suggested by clinical ML governance frameworks (typically >80%), supporting deployment as a clinical decision support tool rather than an autonomous decision-maker.

5.2 External Generalisability

The minimal AUROC degradation (0.016) on the Amsterdam UMCdb external validation -- a European dataset with different patient demographics, clinical protocols, and EHR system compared to the US-based MIMIC-IV -- demonstrates TFT-Health's robust cross-institutional generalisability. XGBoost's larger degradation (0.042) reflects its sensitivity to feature distribution shifts that occur when models are transferred between institutions. TFT-Health's variable selection networks effectively learn which features are reliable across populations and down-weight institution-specific patterns, providing a form of built-in domain adaptation without explicit transfer learning procedures.

5.3 Clinical Implementation Considerations

For clinical deployment, TFT-Health requires integration with EHR data pipelines to extract and normalise the 56 input features in real-time. The model's inference latency (12 ms per prediction on a single GPU) supports real-time alert generation compatible with bedside monitoring systems. Alert fatigue -- the primary barrier to clinical adoption of early warning systems -- can be mitigated by the model's high PPV (42.6% for sepsis at 80% sensitivity), ensuring that most alerts represent genuine clinical deterioration. Integration with clinical workflow should present predictions alongside the attention-highlighted features, enabling clinicians to rapidly assess the clinical plausibility of each alert before acting.

5.4 Limitations

This study has limitations. MIMIC-IV and eICU are US-based datasets; generalisability to low-resource healthcare settings with different disease profiles and EHR completeness is unknown beyond the Amsterdam validation. The prediction tasks use structured EHR data only; incorporating clinical notes (via NLP), imaging, and genomic data may improve predictions but increases complexity. The 48-hour mortality window was chosen for clinical utility but may miss longer-term deterioration patterns. Sepsis labels rely on retrospective Sepsis-3 criteria application, which may not perfectly capture the moment of clinical recognition. The clinician interpretability assessment involved 3 intensivists at a single centre; multi-centre reader studies would provide more generalisable interpretability evidence. Prospective clinical impact studies -- demonstrating improved patient outcomes when TFT-Health alerts are integrated into clinical workflow -- are essential before widespread deployment.

6. Conclusion

This multi-task benchmarking study establishes TFT-Health -- a Temporal Fusion Transformer adapted for clinical EHR prediction -- as the state-of-the-art architecture for temporal healthcare analytics, achieving AUROC 0.894 for 48-hour ICU mortality prediction, 0.762 for 30-day readmission, and 0.882 for 6-hour early sepsis detection, significantly outperforming logistic regression, gradient-boosted trees, random forests, LSTM, and TCN across all tasks. The model's built-in interpretability -- with variable selection networks and temporal attention mechanisms producing clinician-validated explanations (87.4% diagnostic relevance agreement) -- addresses the trust barrier that limits clinical ML adoption. External validation on Amsterdam UMCdb demonstrated robust cross-institutional generalisability with only 0.016 AUROC degradation, the smallest among all models tested. The learned variable importance rankings align with established clinical knowledge, confirming that the model captures genuine physiological deterioration patterns rather than dataset-specific artefacts. Pre-trained TFT-Health models, preprocessing pipelines, and evaluation code are released at <https://github.com/WEDSU-Health/TFT-Health-2023> under MIT licence. Future work should conduct prospective clinical impact trials in ICU settings, extend to multi-modal data (clinical notes, imaging), develop federated learning approaches for privacy-preserving

multi-institutional training, and evaluate TFT-Health as a clinical decision support tool in randomised controlled trials measuring patient-centred outcomes.

References

- Bai, S., Kolter, J.Z. and Koltun, V. (2018) An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271*.
- Calvert, J.S., Price, D.A., Chettipally, U.K., Barton, C.W., Feldman, M.D. and Hoffman, J.L. (2016) A computational approach to early sepsis detection. *Computers in Biology and Medicine*, 74, pp. 69-73.
- Celi, L.A., Mark, R.G., Stone, D.J. and Montgomery, R.A. (2013) Big data in the intensive care unit: closing the data loop. *American Journal of Respiratory and Critical Care Medicine*, 187(11), pp. 1157-1160.
- Finlayson, S.G., Subbaswamy, A., Singh, K., Bowers, J., Kupber, A. and Zittrain, J. (2021) The clinician and dataset shift in artificial intelligence. *The New England Journal of Medicine*, 385(3), pp. 283-286.
- Harutyunyan, H., Khachatrian, H., Kale, D.C., Ver Steeg, G. and Galstyan, A. (2019) Multitask learning and benchmarking with clinical time series data. *Scientific Data*, 6(1), p. 96.
- Johnson, A.E.W., Bulgarelli, L., Shen, L., Gayles, A., Shammout, A. and Horng, S. (2021) MIMIC-IV, a freely accessible electronic health record dataset. *Scientific Data*, 10(1), p. 1.
- Kelly, C.J., Karthikesalingam, A., Suleyman, M., Corrado, G. and King, D. (2019) Key challenges for delivering clinical impact with artificial intelligence. *BMC Medicine*, 17(1), p. 195.
- Kumar, A., Roberts, D., Wood, K.E., Light, B., Parrillo, J.E. and Sharma, S. (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), pp. 1589-1596.
- Li, S., Jin, X., Xuan, Y., Zhou, X., Chen, W. and Wang, Y.X. (2020) Enhancing the locality and breaking the memory bottleneck of transformer on time series forecasting. *Advances in Neural Information Processing Systems*, 32, pp. 5243-5253.
- Lim, B., Arik, S.O., Loeff, N. and Pfister, T. (2021) Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4), pp. 1748-1764.
- Miotto, R., Li, L., Kidd, B.A. and Dudley, J.T. (2016) Deep patient: an unsupervised representation to predict the future of patients from the electronic health records. *Scientific Reports*, 6, p. 26094.
- Rajkomar, A., Oren, E., Chen, K., Dai, A.M., Hajaj, N. and Hardt, M. (2018) Scalable and accurate deep learning with electronic health records. *NPJ Digital Medicine*, 1(1), p. 18.
- Reyna, M.A., Josef, C.S., Jeter, R., Shashikumar, S.P., Westover, M.B. and Nemati, S. (2020) Early prediction of sepsis from clinical data: the PhysioNet/Computing in Cardiology Challenge 2019. *Critical Care Medicine*, 48(2), pp. 210-217.
- Rudd, K.E., Johnson, S.C., Agesa, K.M., Shackelford, K.A., Tsoi, D. and Kievlan, D.R. (2020) Global, regional, and national sepsis incidence and mortality, 1990-2017. *The Lancet*, 395(10219), pp. 200-211.
- Shickel, B., Tighe, P.J., Bihorac, A. and Rashidi, P. (2018) Deep EHR: a survey of recent advances in deep learning techniques for electronic health record analysis. *IEEE Journal of Biomedical and Health Informatics*, 22(5), pp. 1589-1604.
- Tonekaboni, S., Joshi, S., McCradden, M.D. and Goldenberg, A. (2019) What clinicians want: contextualizing explainable machine learning for clinical end use. *Proceedings of Machine Learning Research*, 106, pp. 1-21.
- Zimmerman, J.E., Kramer, A.A., McNair, D.S. and Malila, F.M. (2006) Acute Physiology and Chronic Health Evaluation (APACHE) IV: hospital mortality assessment for today's critically ill patients. *Critical Care Medicine*, 34(5), pp. 1297-1310.
- Zuckerman, R.B., Sheingold, S.H., Orav, E.J., Ruhter, J. and Epstein, A.M. (2016) Readmissions, observation, and the Hospital Readmissions Reduction Program. *The New England Journal of Medicine*, 374(16), pp. 1543-1551.
- Choi, E., Bahadori, M.T., Sun, J., Kulas, J., Schuetz, A. and Stewart, W. (2016) RETAIN: an interpretable predictive model for healthcare using reverse time attention mechanism. *Advances in Neural Information Processing Systems*, 29, pp. 3504-3512.
- Tomassek, A., Clifton, D.A., Georgiou, P., Poon, G. and Shamout, F.E. (2021) Predicting clinical outcomes from large real-world datasets using deep learning. *Nature Biomedical Engineering*, 5(2), pp. 141-153.
- Goldstein, B.A., Navar, A.M., Pencina, M.J. and Ioannidis, J.P.A. (2017) Opportunities and challenges in developing risk prediction models with electronic health records data: a systematic review. *Journal of the American Medical Informatics Association*, 24(1), pp. 198-208.

Declarations

Funding

This study was supported by the Spanish Ministry of Science and Innovation grant PID2022-138642OB-I00 (H. Silva), the French National Research Agency (ANR) grant ANR-22-CE23-0038 (J. Garcia), and the Italian Ministry of University and Research (MUR) grant PRIN 2022RKCWB8 (H. Popescu). Computational

resources were provided by the Spanish Supercomputing Network (RES) under grant BCV-2022-3-0018. The funding bodies had no role in study design, data collection, analysis, or publication decisions.

Conflict of Interest

The authors declare no competing interests. No author holds equity in healthcare AI companies or has received consulting fees from clinical decision support vendors.

Data Availability Statement

MIMIC-IV is available through PhysioNet (physionet.org/content/mimiciv). eICU is available through PhysioNet (physionet.org/content/eicu-crd). Amsterdam UMCdb is available at amsterdammedicaldatascience.nl. All require credentialled access and data use agreements. TFT-Health model code, preprocessing pipelines, and pre-trained weights are available at <https://github.com/WEDSU-Health/TFT-Health-2023> under MIT licence.

Ethical Approval

This study used de-identified, publicly available critical care databases. Access to MIMIC-IV and eICU was granted through PhysioNet credentialling (CITI training completed by all authors). Access to Amsterdam UMCdb was granted under the database's standard data use agreement. No institutional ethics approval was required as all data are de-identified and publicly available for research purposes.

Appendix A

TFT-Health Architecture and Training Configuration

This appendix provides the complete TFT-Health architecture specification, feature preprocessing pipeline, training hyperparameters, and per-task performance breakdown including sensitivity, specificity, and PPV at clinically relevant operating thresholds.

A.1 Architecture Details

- 1a. Variable selection: 2 gated residual networks (GRNs) -- static (14 features -> 128-d) and temporal (42 features x 24 steps -> 128-d per step)
- 1b. GRN structure: FC(input, 128) -> ELU -> FC(128, 128) -> dropout(0.1) -> gating(sigmoid) -> LayerNorm + residual
- 2a. Temporal self-attention: 4 heads, d_model 128, d_k 32; causal mask (attend only to past); interpretable single-head output
- 2b. Position encoding: learned absolute position embeddings (24 time steps) + time-of-day and day-of-week cyclic encoding
- 3a. Multi-task heads: shared encoder -> task-specific FC(128, 64, 1) with sigmoid; mortality and sepsis trained jointly on MIMIC-IV
- 3b. Total parameters: 1.2 M (encoder 0.9 M + heads 0.3 M); inference: 12 ms/prediction (GPU), 45 ms (CPU)

A.2 Training Configuration

- 4a. Optimiser: AdamW (lr 1e-3, betas 0.9/0.98, weight decay 1e-2); cosine annealing with 5-epoch warmup
- 4b. Loss: focal loss (gamma 2, alpha 0.75 for positive class) to address class imbalance; multi-task loss weighted 1:1
- 5a. Epochs: 100 with early stopping (patience 10 on validation AUROC); batch size 256 patients
- 5b. Feature normalisation: z-score per feature using training set statistics; categorical features embedded (dim 8)
- 6a. Missing data: forward-fill within admission, then population median; missingness indicator features appended
- 6b. Data augmentation: random temporal masking (5% of time steps) during training to improve robustness to missing data

A.3 Operating Point Performance

- 7a. Mortality @90% sensitivity: specificity 72.8%, PPV 24.6%, NPV 98.6%, F1 0.381
- 7b. Mortality @80% sensitivity: specificity 82.4%, PPV 32.8%, NPV 97.8%, F1 0.462

- 8a. Sepsis @80% sensitivity: specificity 78.6%, PPV 42.6%, NPV 95.2%, F1 0.548
- 8b. Sepsis @90% sensitivity: specificity 68.2%, PPV 34.8%, NPV 97.4%, F1 0.468