CLUSTER-BASED FEATURE IMPORTANCE LEARNING FOR ELECTRONIC HEALTH RECORD TIME-SERIES

Anonymous authors
Paper under double-blind review

ABSTRACT

The recent availability of Electronic Health Records (EHR) has allowed for the development of algorithms predicting inpatient risk of deterioration and trajectory evolution. However, prediction of disease progression with EHR is challenging since these data are sparse, heterogeneous, multi-dimensional, and multi-modal time-series. As such, clustering is used to identify similar groups within the patient cohort to improve prediction. Current models have shown some success in obtaining cluster representation of patient trajectories, however, they i) fail to obtain clinical interpretability for each cluster, and ii) struggle to learn meaningful cluster numbers in the context of the imbalanced distribution of disease outcomes. We propose a supervised deep learning model to cluster EHR data based on the identification of clinically understandable phenotypes with regard to both outcome prediction and patient trajectory. We introduce novel loss functions to address the problems of class imbalance and cluster collapse, and furthermore propose a feature-time attention mechanism to identify cluster-based phenotype importance across time and feature dimensions. We tested our model in over 100,000 unique trajectories from hospitalised patients with Type-II respiratory failure to predict four different outcomes. Our model yielded added interpretability to cluster formation and outperformed benchmarks by at least 5% in mean AUROC.

1 INTRODUCTION

Chronic conditions such as Chronic Obstructive Pulmonary Disease (COPD) and Cardiovascular Disease (CVD) describe a broad spectrum of medical ailments, and affect a significant percentage of the overall population (Adeloye et al., 2015). Such diseases are characterized by the existence of multiple distinct patient subgroups (corresponding to distinct 'phenotypes'), largely distinguished by differences in pathology and in the response to different treatments and medical interventions Turner et al. (2015); Vogelmeier et al. (2018). Exacerbation of COPD, a condition of respiratory failure, can result in emergency hospital admission and mortality if it is not well treated and managed. Early identification of COPD patients' subgroups is therefore of high medical importance and relevance. EHR time-series data are typically used to determine COPD inpatient phenotypes, and have been applied to detect risk of deterioration (Pikoula et al., 2019).

However, modelling disease progression and risk prediction is challenging due to the extreme data heterogeneity nature of EHRs. Firstly, EHR data contains a mixture of demographic or static variables (i.e. time independent such as age and sex), and multi-dimensional time-series (e.g Heart Rate, HR, and laboratory measurements, such as blood tests). Secondly, EHR time-series are multi-modal as different features are collected from different devices, representing distinct clinical properties of relevance. Similarly, time-series features are sampled at different times and have low and distinct sampling rates, as well as different missing value properties. Furthermore, each feature is associated with different noise and evolution patterns.

Recent advances in deep learning (DL) approaches have shown promising results in EHR modelling due to their capacity to handle complex data (Rajkomar et al., 2018). Nonetheless, DL approaches lack relevant interpretability frameworks to be scaled and applied in hospital settings. Several such models have since been proposed to tackle this issue (Mayhew et al., 2018), however, most of them focus on a subset of EHR features (usually vital signs only) and fail to provide a clinically-focused phenotypic analysis of learnt patient sub-groups (via clustering).

This work builds on previous research in literature to introduce a cluster-based feature-time attention mechanism to predict patient outcomes based on EHR data. Our method also leverages phenotypic information to aid in clinical interpretability, not only making use of demographics and vital-signs information but also of relevant laboratory measurements (all present in the EHR) to provide a more complete patient physiological status. Our contributions include the following:

- An end-to-end DL supervised model to cluster EHR patient data based on the identification of clinically understandable cluster phenotypes with regard to both outcome prediction and patient trajectory in a multi-class setting;
- A weighted loss to address data imbalance for both tasks of clustering and prediction, a common issue in the medical domain;
- The incorporation of a novel loss mechanism in the model, to address the issue of cluster collapse and promote sample assignment to all available clusters;
- Finally, the inclusion of a novel interpretability framework, derived from a cluster-based feature-time attention layer, aiming to identify relevant timestamps and feature variables to represent the patient physiology, cluster assignment and, ultimately, outcome prediction.

This paper is structured as follows. In Section 2, we describe previous research in EHR time-series modelling, clustering and attention methods. Section 3 introduces the dataset used for analysis and provides description of the proposed model. The experimental setup and results of our analysis are presented in Section 4 and discussion takes place in Section 5. Finally, concluding remarks and future work are available in Section 6.

2 Related Work

EHR data comprise complex time-series data, being high-dimensional, multi-modal and heterogeneous, and thus presenting challenges when used in machine learning models (Keogh & Kasetty (2003); Rani & Sikka (2012)). An important goal in a medical setting is to identify phenotypically separable clusters with distinct phenotypic profiles (which we denote as phenotypic clustering hereafter). Phenotypes are clinical variables of interest not used for clustering which provide information about the underlying or future health status. Traditional clustering models such as K-Means or hierarchical clustering have been shown to fail to capture the existing time-dependent feature relationships. As such, variants have been proposed to mitigate this problem. A temporal version of the K-Means algorithm, Time-Series K-Means (TSKM, Tavenard et al. (2020)), models the distance between time-series of different datapoints, using the Euclidean distance (which is equivalent to considering all temporal observations as an independent feature value for the corresponding patient admission), or time-series alignment strategies such as Dynamic-Time Warping (DTW, Berndt & Clifford (1994)) and soft-DTW (Cuturi & Blondel (2017)).

Recent DL architectures, ranging from Auto-Encoders (AE, Ma et al. (2019)), Convolutional Neural Networks (CNN, Munir et al. (2018)) and others, have shown great promise when applied to time-series data across a variety of domains. Fortuin et al. (2019) proposed a Self-Organising Map - Variational Auto-Encoder (SOM-VAE), is a state-of-the-art, unsupervised, DL clustering algorithm which extends a variational auto-encoder architecture (Kingma & Welling, 2013) for observation learning and representation, through the addition of a Markov model (Gagniuc, 2017), to infer temporal evolution within the latent space. Clustering is performed in the low dimensional latent space through the use of self-organising maps (Kohonen (1982)) to obtain a discrete, topologically-interpretable latent representation of the learnt clusters. In a supervised setting, AC-TPC (Lee & Van Der Schaar, 2020) serves as the current state-of-the-art for identifying phenotypically separable clusters of patient trajectory in EHR data. AC-TPC maps EHR data into a latent space via an encoder, and uses an actor-critic network (Konda & Tsitsiklis, 2000) which leverages clinical outcomes to aid in cluster formation and obtaining cluster phenotypes. Neither SOM-VAE and AC-TPC provide clinically meaningful interpretation of feature-time importance or outcome of interest.

Attention mechanisms have recently been proposed to provide greater interpretability to Recurrent Neural Networks (RNN) and to aid in dealing with long-term dependencies (Vaswani et al., 2017; Xu et al., 2015), and have also been used in modelling EHR time-series (Schwab et al., 2017; Shashikumar et al., 2018). RETAIN (Choi et al., 2016) proposes a two-level reverse attention mechanism to

mimic physician's decision process and predict a future diagnosis. In other recent works, attention mechanisms based on bi-directional RNN and CNN outperformed standard classification models in predicting high risk vascular diseases with the addition of medication information as input data (Kim et al., 2017). A drawback of such attention mechanisms is the focus on temporal interpretability only, and inability to look at individual features, which is key in a medical setting. To solve this issue, Shamout et al. (2019) considered independent RNN per feature, with a concatenation of the resulting latent vectors. However, the latter does not allow the joint modelling across both feature and time dimensions. Alternatively, Kaji et al. (2019); Gandin et al. (2021) proposed learning attention weights directly on the original inputs, prior to being transformed by a RNN, which does not allow modelling of the resulting latent representations. To the best of our knowledge, no existing models have been proposed that jointly leverage both feature and time dimensions (feature-time) to determine clinical observation relevance on clustered EHR data.

3 METHODS

3.1 DATASET AND PRE-PROCESSING

Our dataset was retrieved from a retrospective database of routinely collected observations from concluded hospital admissions between March 2014 and March 2018 (the HAVEN project, REC reference: 16/SC/0264 and Confidential Advisory Group reference 08/02/1394). The database includes EHR measurements of adult patients admitted to four hospitals from the Oxford University Hospitals NHS Foundation Trust. Note that the HAVEN dataset does not include data from Intensive Care Units (ICU), and we have excluded observations taken in the Emergency Department. Key characteristics of HAVEN cohort data include a) heterogeneity, b) multi-modality, and difference in: c) noise distributions, d) sampling rates, e) missing values, etc. Such properties are common across EHR settings, and are challenging with respect to learning useful representations and predictions.

We used the protocol defined in Pimentel et al. (2019) to subset the cohort to those patients at risk of developing Type-II Respiratory Failure (T2RF) in hospital (a diagram of the data selection steps can be found in Figure A.1 in the Appendix). Four patient outcomes were considered in our analysis: i) no event during hospital stay, leading to successful discharge from the hospital, or the first instance of one of three possible events, ii) unplanned entry to ICU, iii) cardiac arrest (also named "Cardiac" hereafter) and iv) death. Outcome groups are not clearly separable (see Tables A.2, A.3 in the Appendix), so patient clusters will naturally contain a mix of different admission outcomes. In this setting, a cluster phenotype is a learnable variable represented as a categorical distribution indicating the corresponding propensity for cluster-assigned patients to each corresponding outcome.

For each admission, observations were grouped according to mean window observation value into 4 hour blocks based on the time to outcome (discharge in the case of no event during stay) - only observations within 24 and 72 hours before the outcome were considered. This time window was selected based on those traditionally used for validating Early Warning Score (EWS) systems (baseline models used by UK NHS staff to track inpatient physiology, (Royal College of Physicians, 2017) and clinical input, such that the target phenotype represents the patient status in the subsequent 24 hours. Features were transformed according to min-max normalization due to skewness and heterogeneity in their distributions. Patient admissions were randomly split into train, validation and test sets. Missing values were imputed based on the previously observed time block - all remaining missing observations were imputed according to the feature median from the aggregated validation and test data (see Section 4 for description of train-test data split). Imputed values were flagged in a three-dimensional mask matrix.

After processing, input data contained over 100,000 unique patient trajectories corresponding to 4,266 unique patient admissions (only last patients admissions were considered in our analysis). Original trajectories for the patient cohort are shown in the Appendix in Figures A.4, A.5, A.6 for different variables/features. A lack of clear outcome group separability can be observed across temporal and static variables. Furthermore, we note the high degree of imbalance in the data - admissions with no event account for over 86.8% of the total number of admissions, while event classes correspond to 10.3% Death, 1.8% ICU and 1.1% Cardiac.

3.2 PROPOSED MODEL

We propose a novel model, which we denote by **C**luster-b**A**sed i**M**portanc**E** **L**earning f**O**r **T**ime-series (CAMELOT). Our proposed methodology is displayed in Figure 1[1]. Our model builds on previous literature on 3 key items: a) a modified loss function to target the multi-class imbalance, b) a novel loss function to ensure cluster assignment and phenotype are representative, and c) a novel feature-time attention-level framework to boost representation and introduce feature-time interpretability for cluster assignment.

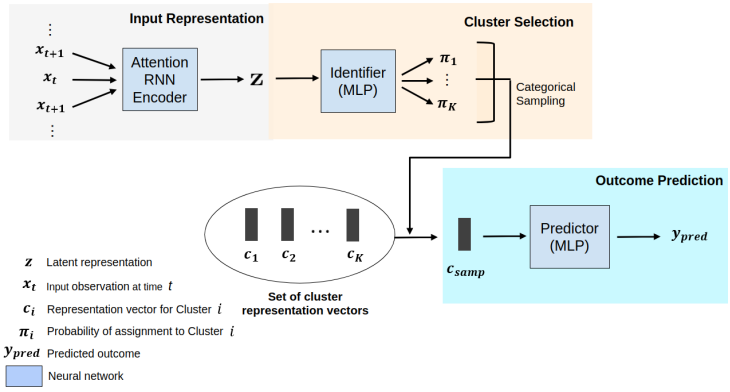

Figure 1: Diagram of proposed model. MLP - Multilayer Perceptron neural network blocks; RNN - Recurrent Neural Network.

Let N denote the number of patients and D_f the number of features. Input data consists of a set of patient trajectories $\mathbb{X} = \{\{\mathbf{x}_{n,t}\}_{t=1}^{T_n}\}_{n=1}^{N}$, where T_n is the maximum number of temporal observations for patient n, and a set of patient outcomes $\mathbb{Y} = \{\mathbf{y}_n\}_{n=1}^{N}$. Input trajectory data for the n-th patient is represented as $\mathbf{X}_n = [\mathbf{x}_{n,1}, ..., \mathbf{x}_{n,T_n}]$, where each $\mathbf{x}_{n,t} \in \mathbb{R}^{D_f}$ is referred to as an *observation* (vector), with a maximum of observed D_f feature values. The corresponding patient outcome is a *one-hot encoded* vector $\mathbf{y}_n \in \mathbb{R}^4$ (more generally, the dimension of $\mathbf{y}_n$ equals to the number of possible outcomes).

Our DL model can be decomposed into 3 neural network blocks: an Encoder, Identifier and Predictor. We refer the action of each network respectively as E, I and P (for example, $I(\mathbf{x})$ denotes the output of the Identifier given some input vector $\mathbf{x}$). Both the Identifier and Predictor are Multilayer Perceptrons (MLP), networks of stacked feed-forward dense layers. On the other hand, the Encoder block can be further sub-divided into a) a stack of RNN layers and b) our proposed custom attention layer (see Section 3.3 for further details). Separately, we also consider a set of *trainable* cluster representation vectors, $\mathcal{C} = \{\mathbf{c}_1, ..., \mathbf{c}_K\}$. We assign the phenotype for cluster i as $P(\boldsymbol{c}_i)$.

A model call is as follows: Given the n-th patient input trajectory data $\mathbf{X}_n$, the Encoder network returns a *latent representation* $\mathbf{z}_n := E(\mathbf{X}_n) \in \mathbb{R}^l$. Consequently, the Identifier network computes cluster assignment probabilities, $\boldsymbol{\pi}_n := I(\mathbf{z}_n) \in \mathbb{R}^K$. Each element of $\boldsymbol{\pi}_n$, π_n^i, represents the probability assignment of $\mathbf{z}_n$ to cluster i, given a total of K clusters. A cluster, k_{samp}^n is selected according to categorical sampling ($k_{\text{samp}}^n \sim \text{Cat}(\boldsymbol{\pi}_n)$), and the corresponding cluster representation, $\mathbf{c}_{\text{samp}}^{\mathbf{n}} := \mathbf{c}_{\mathbf{k}_{\text{samp}}^{\mathbf{n}}}$ is then selected from $\mathcal{C}$. The output of the model is $y_{\text{pred}} := P(\mathbf{c}_{\text{samp}}^n) \in \mathbb{R}^4$. We note that k_{samp}^n is only sampled during a *training phase*; at prediction stage, cluster selection follows the equation $k_{\text{pred}}^n = \underset{i=1,...,K}{\arg\max}\, \pi_n^i$.

3.3 ENCODER NETWORK AND A CUSTOM ATTENTION LAYER

The diagram of our proposed Encoder network is presented in Figure 2. The Encoder contains (i) a Recurrent Neural Network (RNN) block of stacked Long Short-Term Memory (LSTM) layers, and (ii) a customised attention layer, which computes a latent representation by comparing input

[1]code will be shared online after review

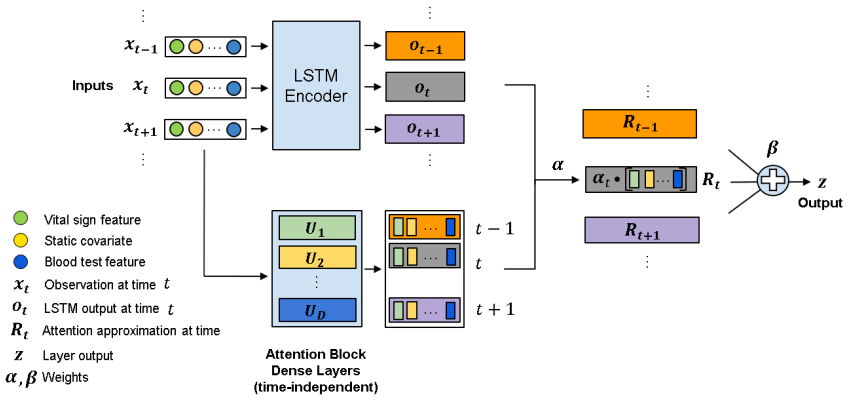

Figure 2: Diagram of the Encoder network composed of an LSTM Encoder and a custom attention layer.

data with the sequence of output states from the RNN block. We use the same notation as above, and write the sequence of output states of the final LSTM layer as $\mathbf{o}_{n,1}, ..., \mathbf{o}_{n,T_n}$, with $\mathbf{o}_{n,i} \in \mathbb{R}^l$. Theoretically, each $\mathbf{o}_{n,t}$ corresponds to a representative summary of input patient information up until time t. We propose to approximate $\mathbf{o}_{n,t}$ as a linear combination of latent representations of each individual feature, thereby allowing the separation of output states into contributions from each feature. Note that it is important feature transformations be *time-independent*, to avoid over-parametrising the model, over-fitting and ensure feature representation maps are similar across time.

Our attention layer behaves as a set of D_f feed-forward neural network layers, $\mathbf{U}_1, ..., \mathbf{U}_{D_f}$, jointly represented by: (i) a matrix of *learnable* kernel weights $\boldsymbol{D} \in \mathbb{R}_{l \times D_f}$. We write $\boldsymbol{D} = [\mathbf{D}_1, ..., \mathbf{D}_{D_f}]$; (ii) a matrix of *learnable* bias vectors $\boldsymbol{B} \in \mathbb{R}_{l \times D_f}$. Similarly, we can write $\boldsymbol{B} = [\mathbf{B}_1, ..., \mathbf{B}_{D_f}]$; and (iii) an activation function, σ which matches the output activation of the RNN block.

Input data for patient n, $\mathbf{X}_n$ is fed as input to the RNN block, which outputs a sequence of latent output states $(\mathbf{o}_{n,t})_{t=1}^{T_n}$. For $t = 1, ..., T_n$, we compute D_f feature representations in latent space as:

$$\boldsymbol{R}_{n,t} := \sigma(\boldsymbol{D} \odot \mathbf{x}_{n,t} + \boldsymbol{B}) \tag{1}$$

where $\boldsymbol{R}_{n,t} = [\mathbf{R}_{n,t}^1, ..., \mathbf{R}_{n,t}^{D_f}]$ is our collection of feature representations, σ is applied element-wise and $\mathbf{A} := \boldsymbol{D} \odot \mathbf{x}_{n,t}$ is a matrix satisfying $A_{i,j} = \boldsymbol{D}_{i,j}(\mathbf{x}_{n,t})_j$. Equivalently, $\mathbf{R}_{n,t}^i$ is the output of a dense layer, $\mathbf{U}_i$ with kernel $\mathbf{D}_i$, bias $\mathbf{B}_i$, activation σ and input $(\mathbf{x_{n,t}})_i$. We approximate $\mathbf{o}_{n,t} \approx \sum_{i=1}^{D_f} \alpha_t^i \mathbf{R}_{n,t}^i = \boldsymbol{R}_{n,t}\boldsymbol{\alpha}_t$. This approximation is minimised following a least squares criterion, which can has a well-known solution, $\widehat{\boldsymbol{\alpha}}_t$, and corresponding optimal approximation $\widehat{\mathbf{o}}_{n,t}$.

Given, $\hat{\mathbf{o}}_{n,t}$, a similar procedure is used to compute a context vector as

$$\mathbf{z} := \sum_t \beta_t \hat{\mathbf{o}}_{n,t} \tag{2}$$

where weights $\boldsymbol{\beta}$ are learned to provide a more representative context vector.

3.4 Attention Map Visualisation

Given cluster representation vectors, $\mathbf{c}_k$, we can compute a cluster-wise feature-time attention visualisation map as follows. First, we normalise feature-weights $\widehat{\boldsymbol{\alpha}}_t$ according to a softmax function, $\mathbf{s}_t = \sigma(\widehat{\boldsymbol{\alpha}}_t) \in \mathbb{R}^{D_f}$, where σ is the softmax function:

$$\sigma(\mathbf{x}) = \frac{\exp|\mathbf{x}|}{\|\exp|\mathbf{x}|\|_1} \tag{3}$$

Secondly, we compute cluster-wise weights, γ_t^k according to a least-square approximation of $\mathbf{c}_k \approx \sum_{t=1}^{T_n} \widehat{\mathbf{o}}_{n,t}\gamma_{n,t}^k$, and solved as before. We similarly normalise $\gamma_{n,t}^k$ to obtain cluster temporal scores, $e_{n,t}^k = \sigma(\gamma_{n,t}^k)$. Finally, we can compute K scoring matrices, $M_n^1, ..., M_n^K \in \mathbb{R}_{T_n \times D_f}$:

$$\left(M_n^k\right)_{t,f} = e_{n,t}^k \boldsymbol{s}_t^f \tag{4}$$

Note that:

$$\|M_n^k\|_1 = \sum_t e_{n,t}^k \sum_f \boldsymbol{s}_t^f = \sum_t e_{n,t}^k = 1$$

Given that Matrices M_n^k are normalised, they may be consequently, visualised as a normalised feature-time map for cluster assignment relevance and provide further model interpretability.

3.5 Loss optimisation

The model is optimised through consideration of three distinct loss functions. We introduce a weighted cross-entropy loss function:

$$L_{\text{pred}}(y_{\text{true}}, y_{\text{pred}}) = -\sum_{c=1}^{C} w_c y_{\text{true}}^c \log\left(y_{\text{pred}}^c\right) = -w_{c'} \log(y_{\text{pred}})_{c'} \tag{5}$$

where c' is the true outcome for a particular patient. We propose inversely proportional normalised weights: $\sum_{c=1}^{C} w_c = 1$ and w_c is inversely proportional to the class distribution, i.e., $w_c \propto \frac{N}{N_c}$ with N being the number of patients and N_c being the number of patients with outcome label c. Class weighting penalises misclassification more heavily on less sampled classes.

We also propose a novel distribution loss function, $L_{\text{dist}}(\boldsymbol{\pi})$. We define the average cluster probability of assignment as $\boldsymbol{\pi}_C := \frac{1}{N}\sum_n \boldsymbol{\pi}_n$. Then we introduce $L_{\text{dist}}(\boldsymbol{\pi}) = -H(\boldsymbol{\pi}_C)$, where H denotes entropy. Note that L_{dist} is minimised when π_C is uniform, ensuring all clusters are 'explored' and have comparable number of samples. While clusters should not necessarily be assigned the same number of samples, this loss helps to overcome a concern of *cluster collapse*, where clusters do not separate and samples are assigned to a very small number of non-representative clusters. Finally, to separate cluster representation vectors, we define the cluster separation loss as $L_{\text{clus}}(\mathcal{C}) = -\frac{1}{K(K-1)}\sum_{i,j}\|\mathbf{c}_i - \mathbf{c}_j\|^2$

To optimise our model, iterative gradients are applied according to weighted combinations of the above loss functions with hyper-parameter weights α, β:

1. Outcome Predictor is updated according to L_{pred};
2. Encoder and Identifier are trained according to $L_{\text{pred}} + \alpha L_{\text{dist}}$;
3. Cluster representation vectors are updated with regards to $L_{\text{pred}} + \beta L_{\text{clus}}$.

3.6 Initialisation

Our proposed model also follows a set of initialisation pre-training procedures. Firstly, the Encoder and Outcome Predictor are pre-trained according to a classification task ($\tilde{\mathbf{y}} = P(E(\mathbf{x}))$), with corresponding loss L_{pred}. Latent state representations $E(\mathbf{x})$ are clustered through a K-means algorithm with K clusters across the whole training set. Cluster representation vectors are initialised as given by the resulting cluster centroids, and finally the Cluster Identifier network is pre-trained to identify clusters as predicted by the K-Means algorithm with categorical cross-entropy loss. Implementation was completed in Python, with TensorFlow 2, scikit-learn and NumPy. All experiments were run with 1 Tesla v100 GPU, and 8 CPUs Intel(R) Xeon(R) Gold 6246 @ 3.30GHz.

4 Results

For Benchmark purposes, we considered TSKM as a classic clustering benchmark, and SOM-VAE and AC-TPC as state-of-the-art phenotypic clustering methods. AC-TPC considers temporal subsequences of a complete patient set of observations - for comparison purposes, we consider only the model output for the complete patient sequence. For simplicity, we present results with all input features considered.

All models were trained on the same training set (60% of the complete input data) and evaluated against the same test set (remaining 40%). For DL models, we further split the training set into a purely training and validation sets. All experiments with varying hyper-parameters were repeated 10 times with a fixed set of 10 distinct seeds, and results are reported according to average metric and

standard deviation. A complete list of the hyper-parameters considered for each model is included in Table A7 in the Appendix. In bold, top-performing hyper-parameters are indicated. Optimal integer hyper-parameters (K, l) were selected according to an "Occam's Razor" approach - for each parameter, we assign it the highest value such that increasing this amount does not lead to a significant increase in performance according to mean AUROC and a Friedman's hypothesis test. Neural network size parameters were kept consistent across all DL models where applicable. All other optimal hyper-parameters were selected according to highest AUROC performance conditional on the model predicting each class (e.g. not making predictions solely for no event or Death Events).

We evaluated clustering performance through standard clustering metrics, including Silhouette score (SIL, Rousseeuw (1987)), Davies-Bouldin Index (DBI, Davies (1979)), Variance Ratio Criterion (VRI, Caliński & Harabasz (1974)). Results for all clustering models are displayed in Table 3. In Table 4, we evaluated the (multi-class) prediction performance with regards to Area-under-the-Receiver-Operating-Curve (AUROC), unweighted mean F1-score, unweighted mean Recall, and Normalised Mutual Information (NMI). For purely unsupervised models (SOM-VAE and TSKM), an outcome predictive pipeline was constructed by assigning patient admissions to clusters, and consequently to the empirical outcome distribution in the corresponding cluster. The prediction task was also benchmarked against traditional classifiers for outcome prediction in EHR data, namely Support Vector Machines (SVM), XGBoost (XGB) and NEWS2 (i.e., the National Early Warning Score used in the UK hospitals).

Metric	TSKM	SOM-VAE	AC-TPC	CAMELOT (proposed)
SIL	**0.35** (±**0.01**)	0.25 (± 0.08)	0.04 (± 0.01)	0.11 (±0.04)
DBI	**1.19** (±**0.08**)	1.89 (± 0.63)	4.34 (± 0.80)	3.12 (±0.53)
VRI	**554.6** (±**2.50**)	12.8 (± 9.32)	66.5 (± 18.7)	216.7 (±6.2)

Table 3: Clustering separability results by the different clustering methodologies given input data with all available features (static, vital-signs, serum and haematological variables). For each metric and model, the average score and standard deviation are returned. The best values for each metric are indicated in bold.

Metric	AUROC	F1-score	Recall	NMI
SVM	0.50 (± 0.02)	0.23 (± 0.00)	0.25 (± 0.00)	0.01 (± 0.02)
XGB	0.65 (± 0.01)	0.23 (± 0.00)	0.22 (± 0.00)	0.03 (± 0.04)
NEWS2	0.61	0.29	0.34	0.01
TSKM	0.55 (± 0.01)	0.24 (± 0.03)	0.26 (± 0.02)	0.01 (± 0.03)
SOM-VAE	0.61 (± 0.09)	0.27 (± 0.05)	0.27 (± 0.03)	0.05 (± 0.03)
AC-TPC	0.68 (± 0.01)	**0.38** (±**0.01**)	0.36 (± 0.01)	0.17 (± 0.02)
CAMELOT (proposed)	**0.73** (±**0.02**)	0.36 (±0.01)	**0.38** (±**0.02**)	**0.20** (±**0.03**)

Table 4: Outcome prediction scores across all models, displayed with an average and standard deviation of a set of 10 seeds (except NEWS2, which is deterministic). The best values for each metric are indicated in bold.

On top of performance evaluation with regards to clustering separability and outcome prediction, we display below a comparison between the learnt cluster phenotypes of the proposed model and that of phenotypic clustering benchmark AC-TPC. For each cluster, the corresponding phenotype $P(\mathbf{c})$ is shown as a bar plot over the 4 possible outcomes with corresponding phenotypic probability value. We also display cluster phenotype plots for both TSKM and SOM-VAE in the Appendix (Figures A.8 and A.9). We note that the cluster phenotypes learnt by CAMELOT also align with the empirical outcome relevance in the learnt clusters (Table A.10). Lastly, we also display feature-time cluster relevance attention maps in Figure 6. For each cluster, a patient was randomly selected from the set of patients in the corresponding cluster, and a corresponding feature-time attention matrix and visualised as a heatmap.

5 Discussion

Our proposed model shows an improvement in clustering performance (see Table 3) when compared to the current phenotypic clustering benchmark (AC-TPC), and outperforms SOM-VAE according to VRI. Although the cluster separability metrics are superior in the case of TSKM, this is expected given a metric bias towards convex clusters, and the convexity of the K-Means based algorithm. Furthermore, DL clustering occurs in a latent space, which, unfortunately, is not directly comparable with separability results on the input data.

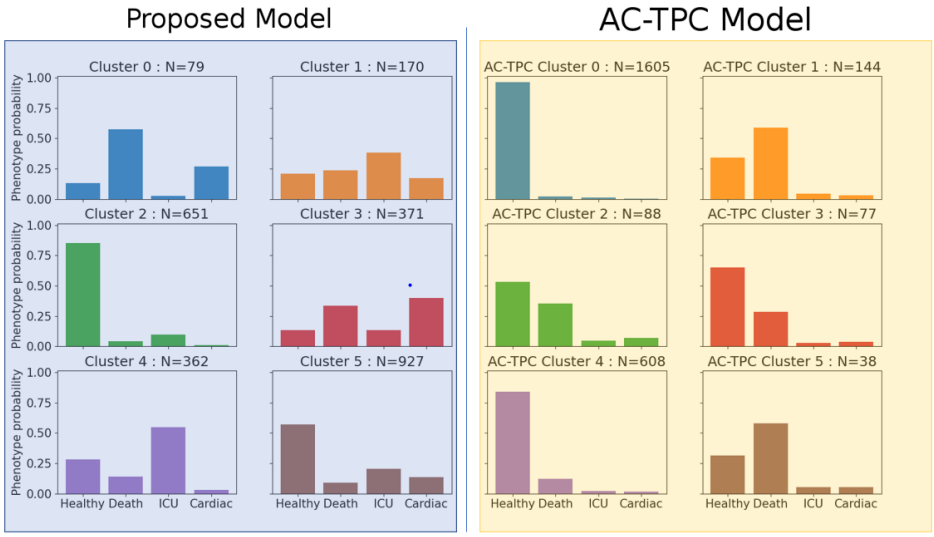

Figure 5: Comparison of bar plots of cluster phenotypes for the proposed model and benchmark AC-TPC. On the left (blue), phenotypes are displayed for each cluster (out of a total of 6), and each phenotype corresponds to the probability of an outcome. Similar results shown on the right (yellow) for AC-TPC. The title of each sub-plot indicates the cluster considered, as well as the number of patients assigned to a given cluster.

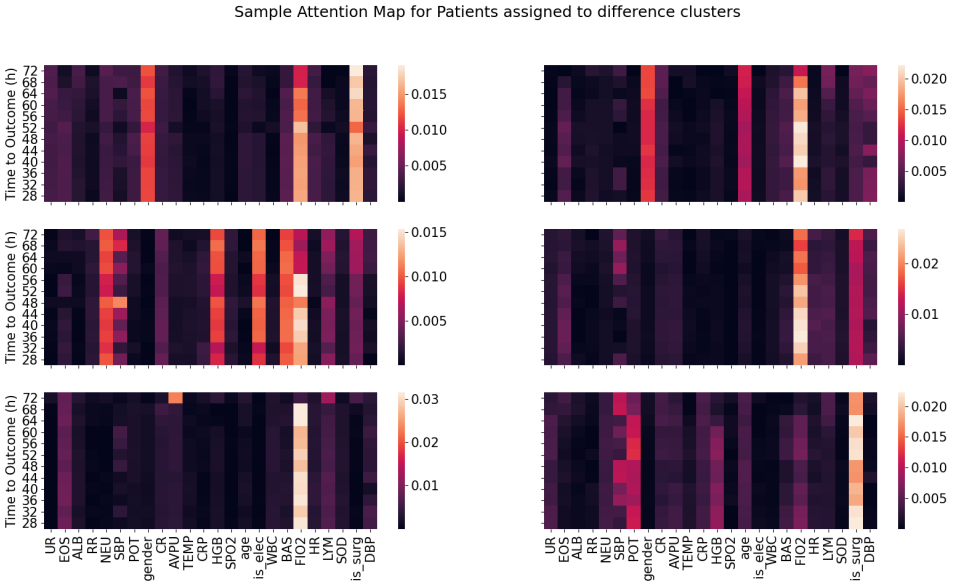

Figure 6: Feature-Time Cluster Relevance Map. Each Heatmap represents a feature-time relevance matrix for a random patient assigned to a given cluster. Vertical Axis indicates time to outcome, in hours, while horizontal axis indicates different input features.

With regards to predictive power, it can be seen in Table 4 that our model outperforms both standard classifiers (at least 8%) and benchmarking clustering methods (at least 5%) according to mean AUROC. A similar increase can be seen in other classification metrics, with the exception of F1-score, where model performance is slightly below to that of AC-TPC. Our model is able to more accurately determine patterns in the data than the previously proposed models, as EHR data is extremely complex and heterogeneous. It is particularly promising that the model obtains good predictive task results despite a clustering bottleneck (i.e. sample predicted outcomes are done through the assigned cluster, as opposed to tailored to the precise input data). While it is possible that other models could

show better performance on the direct task of outcome prediction given EHR input, such models can potentially be associated with lack of robustness or input sensitivity difficulties. Furthermore, it is likely they would struggle with identifying relevant trends and properties of clinical interest. As such, for new admissions, such models could provide a prediction for the overall outcome, but no robust understanding of how this outcome will occur, and how to prevent potential risks of deterioration, let alone the ability to pool data from other similar patients.

Figure 5 shows the advantage of two key aspects of our methodology. The identification of clear, separable cluster phenotypes provides a useful layer of interpretability to clinicians as it characterises the profile and status representation for each cluster according to chances of potential different outcomes. Our model identifies clearer and more separable cluster phenotypes than AC-TPC (other benchmark results in Appendix A.8 and A.9), and allows a clearer look into the underlying patient status for the corresponding clusters. Clusters 0 and 3 are largely representative of death and cardiac events on the subsequent 24 hours, while cluster $1, 4, 5$ are healthier, with a smaller change of death and/or a cardiac arrest. Cluster 2 is a "healthy" cluster. On the other hand, phenotypes learnt by AC-TPC are unable to provide this detail of information, in particular, without the existence of a single clear and unique cluster phenotype. See Appendix A.10 for the true number of patients per outcome for each cluster. Note, furthermore, that the model managed to successfully navigate a heavy class-imbalance setting. Representative clusters are able to capture different-sized sub-populations (as can be seen by the difference in cluster sizes), yet still identify the potential risks of deterioration.

On the other hand, learnt cluster attention maps introduced in Figure 6 introduce yet another layer of interpretability to our proposed clustering model. By separating latent space representation into components derived from each feature, we are able to determine, for each patient, the key feature-time observations of relevance for the corresponding patient. Attention maps highlight: a) common feature-time regions of interest across clusters, which can indicate general key regions for useful patient physiology representation and b) cluster-unique feature-time areas. For instance, less healthy clusters (say, Clusters 0 and 5 in Figure 5), both present different unique zones of high relevance - potential useful information for clinicians to further understand the cluster-defined sub-population, as well as potential deterioration signals in a clinical ward.

6 Conclusion and Future Work

In this work, we propose a novel deep learning model for the task of identification of phenotypically separable clusters applied to EHR data for. As part of our proposed model, we propose 2 distinct loss functions and introduce a novel feature-time attention layer to better represent patient data and to introduce a feature-time relevance map for each cluster. Our experiments show promising results with the addition of both methodological tools above, on both cluster separability and outcome prediction performance. The addition of the feature-time layer has the added benefit of introducing key interpretability tools for researchers to understand relevant regions for good patient physiology representation as well as an indication of what can lead to patient deterioration.

There are multiple interesting avenues of investigation building on this work. On the other hand, the current attention layer mechanisms could potentially be improved with the addition of temporal weight smoothness, or, alternatively, weight regularization to encourage exploration of the complete feature-time space. Alternatively, cluster selection through a neural network mechanism introduces high capacity at the potential cost of robustness and cluster collapse. Potentially, more traditional methods incorporated into a similarly complex pipeline can achieve better performing through a clearer identification of cluster regions in latent space. Furthermore, methodological improvements will also benefit from a more extensive testing across other diverse datasets and other potential areas of application.

References

Davies Adeloye, Stephen Chua, Chinwei Lee, Catriona Basquill, Angeliki Papana, Evropi Theodoratou, Harish Nair, Danijela Gasevic, Devi Sridhar, Harry Campbell, et al. Global and regional estimates of copd prevalence: Systematic review and meta–analysis. *Journal of global health*, 5 (2), 2015.

Donald J Berndt and James Clifford. Using dynamic time warping to find patterns in time series. In *KDD workshop*, volume 10, pp. 359–370. Seattle, WA, USA:, 1994.

Tadeusz Caliński and Jerzy Harabasz. A dendrite method for cluster analysis. *Communications in Statistics-theory and Methods*, 3(1):1–27, 1974.

Edward Choi, Mohammad Taha Bahadori, Joshua A Kulas, Andy Schuetz, Walter F Stewart, and Jimeng Sun. Retain: An interpretable predictive model for healthcare using reverse time attention mechanism. *arXiv preprint arXiv:1608.05745*, 2016.

Marco Cuturi and Mathieu Blondel. Soft-dtw: a differentiable loss function for time-series. In *International Conference on Machine Learning*, pp. 894–903. PMLR, 2017.

DL Davies. et dw bouldin. a cluster separation measure. *IEEE Trans. Pattern Anal. Mach. Intell*, 1 (2), 1979.

Vincent Fortuin, Matthias Hüser, Francesco Locatello, Heiko Strathmann, and Gunnar Rätsch. Deep self-organization: Interpretable discrete representation learning on time series. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=rygjcsR9Y7.

Paul A Gagniuc. *Markov chains: from theory to implementation and experimentation*. John Wiley & Sons, 2017.

Ilaria Gandin, Arjuna Scagnetto, Simona Romani, and Giulia Barbati. Interpretability of time-series deep learning models: A study in cardiovascular patients admitted to intensive care unit. *Journal of Biomedical Informatics*, 121:103876, 2021. ISSN 1532-0464. doi: https://doi.org/10.1016/j.jbi.2021.103876. URL https://www.sciencedirect.com/science/article/pii/S1532046421002057.

Deepak A Kaji, John R Zech, Jun S Kim, Samuel K Cho, Neha S Dangayach, Anthony B Costa, and Eric K Oermann. An attention based deep learning model of clinical events in the intensive care unit. *PloS one*, 14(2):e0211057, 2019.

Eamonn Keogh and Shruti Kasetty. On the need for time series data mining benchmarks: a survey and empirical demonstration. *Data Mining and knowledge discovery*, 7(4):349–371, 2003.

You Jin Kim, Yun-Geun Lee, Jeong Whun Kim, Jin Joo Park, Borim Ryu, and Jung-Woo Ha. Highrisk prediction from electronic medical records via deep attention networks. *arXiv preprint arXiv:1712.00010*, 2017.

Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

Teuvo Kohonen. Self-organized formation of topologically correct feature maps. *Biological cybernetics*, 43(1):59–69, 1982.

Vijay R Konda and John N Tsitsiklis. Actor-critic algorithms. In *Advances in neural information processing systems*, pp. 1008–1014, 2000.

Changhee Lee and Mihaela Van Der Schaar. Temporal phenotyping using deep predictive clustering of disease progression. In Hal Daumé III and Aarti Singh (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 5767–5777. PMLR, 13–18 Jul 2020. URL http://proceedings.mlr.press/v119/lee20h.html.

Qianli Ma, Jiawei Zheng, Sen Li, and Gary W Cottrell. Learning representations for time series clustering. *Advances in neural information processing systems*, 32:3781–3791, 2019.

Michael B Mayhew, Brenden K Petersen, Ana Paula Sales, John D Greene, Vincent X Liu, and Todd S Wasson. Flexible, cluster-based analysis of the electronic medical record of sepsis with composite mixture models. *Journal of biomedical informatics*, 78:33–42, 2018.

Mohsin Munir, Shoaib Ahmed Siddiqui, Andreas Dengel, and Sheraz Ahmed. Deepant: A deep learning approach for unsupervised anomaly detection in time series. *Ieee Access*, 7:1991–2005, 2018.

Maria Pikoula, Jennifer Kathleen Quint, Francis Nissen, Harry Hemingway, Liam Smeeth, and Spiros Denaxas. Identifying clinically important copd sub-types using data-driven approaches in primary care population based electronic health records. *BMC medical informatics and decision making*, 19(1):1–14, 2019.

Marco AF Pimentel, Oliver C Redfern, Stephen Gerry, Gary S Collins, James Malycha, David Prytherch, Paul E Schmidt, Gary B Smith, and Peter J Watkinson. A comparison of the ability of the national early warning score and the national early warning score 2 to identify patients at risk of in-hospital mortality: a multi-centre database study. *Resuscitation*, 134:147–156, 2019.

Alvin Rajkomar, Eyal Oren, Kai Chen, Andrew M Dai, Nissan Hajaj, Michaela Hardt, Peter J Liu, Xiaobing Liu, Jake Marcus, Mimi Sun, et al. Scalable and accurate deep learning with electronic health records. *NPJ Digital Medicine*, 1(1):1–10, 2018.

Sangeeta Rani and Geeta Sikka. Recent techniques of clustering of time series data: a survey. *International Journal of Computer Applications*, 52(15), 2012.

Peter J Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65, 1987.

Royal College of Physicians. National early warning score (news) 2. *Standardising the assessment of acute-illness severity in the NHS*, 2017.

Patrick Schwab, Gaetano C Scebba, Jia Zhang, Marco Delai, and Walter Karlen. Beat by beat: Classifying cardiac arrhythmias with recurrent neural networks. In *2017 Computing in Cardiology (CinC)*, pp. 1–4. IEEE, 2017.

Farah E Shamout, Tingting Zhu, Pulkit Sharma, Peter J Watkinson, and David A Clifton. Deep interpretable early warning system for the detection of clinical deterioration. *IEEE journal of biomedical and health informatics*, 24(2):437–446, 2019.

Supreeth P Shashikumar, Amit J Shah, Gari D Clifford, and Shamim Nemati. Detection of paroxysmal atrial fibrillation using attention-based bidirectional recurrent neural networks. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 715–723, 2018.

Romain Tavenard, Johann Faouzi, Gilles Vandewiele, Felix Divo, Guillaume Androz, Chester Holtz, Marie Payne, Roman Yurchak, Marc Rußwurm, Kushal Kolar, et al. Tslearn, a machine learning toolkit for time series data. *J. Mach. Learn. Res.*, 21(118):1–6, 2020.

Alice M Turner, Lilla Tamasi, Florence Schleich, Mehmet Hoxha, Ildiko Horvath, Renaud Louis, and Neil Barnes. Clinically relevant subgroups in copd and asthma. *European respiratory review*, 24(136):283–298, 2015.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.

Claus F Vogelmeier, Kenneth R Chapman, Marc Miravitlles, Nicolas Roche, Jørgen Vestbo, Chau Thach, Donald Banerji, Robert Fogel, Francesco Patalano, Petter Olsson, et al. Exacerbation heterogeneity in copd: subgroup analyses from the flame study. *International journal of chronic obstructive pulmonary disease*, 13:1125, 2018.

Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pp. 2048–2057. PMLR, 2015.

A APPENDIX

A.1 DATA

A description of the complete pipeline of data re-processing, following the protocol defined in Pimentel et al. (2019), is shown in Figure A.1.

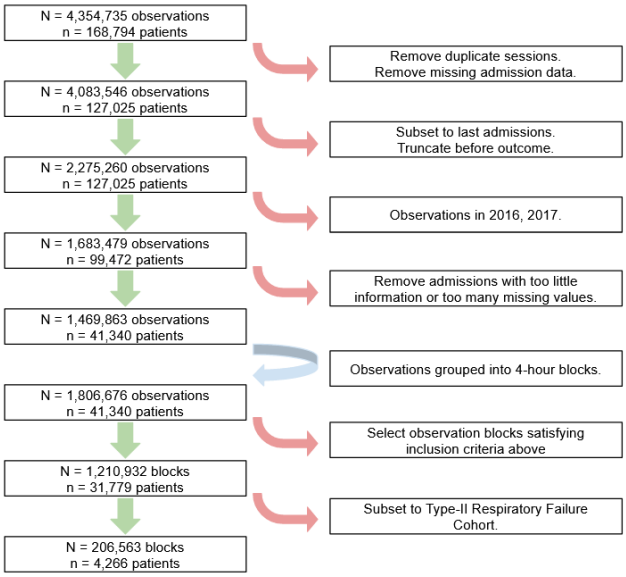

A.1: HAVEN processing

A total of 26 input features were considered. Firstly, 4-hourly vital-sign sets which included 8 features: Heart Rate (HR), Respiratory Rate (RR), Systolic Blood Pressure (SBP), Diastolic Blood Pressure (DBP), peripheral Oxygen Saturation (SpO_2), Temperature (TEMP), level of consciousness via the AVPU scale - Alert, Verbal, Pain, Unresponsive - and estimated Fraction of Inspired Oxygen (FiO_2, available when an oxygen mask is applied to the patient). Each set consisted of a timestamp and the vital-sign numerical values. Secondly, 4 demographic variables were selected (modelled as static variables): age, sex, and admission type (elective or surgical). Thirdly, we included 6 features resulting from biochemistry blood tests, denoted as 'Serum': Serum levels of urea, albumin, creatinine, sodium, potassium and C-reactive protein. Finally, 8 haematological blood test features were also included: white and haemaglobin cell counts, concentration of eosinophils, basophils, neutrophils, and lymphocytes, as well as eosinophil-to-basophil and neutrophil-to-lymphocyte ratios. These features were selected based on domain knowledge of features related to severity in the prognosis and outcome of T2RF inpatients.

Descriptive statistics for all input variables is described in Table A.2. Median and inter-quartile range (IQR) is displayed for continuous and categorical variables, while binary variables are shown according to number of counts in the dataset and corresponding cohort proportion. Statistics are displayed for the complete data ("All"), but also for each sub-cohort defined by the overall outcome. We can observe that these sub-cohorts are not clearly separable and are hard to identify solely from this information.

A summary of the patient cohort in relation to outcomes and target phenotypes can be seen in Table A.3. Challenges with regards to obtaining phenotypically separable clusters can similarly be observed - there is no clear significant difference between the target outcome sub-cohorts with regards to demographic input variables. With regards to outcome distribution, we also note the high degree of imbalance in the dataset - the large majority of the patients in our dataset suffered from no adverse events (over 86%), while only 48 had a Cardiac event, and 76 were re-directed to the ICU.

The lack of outcome sub-cohort separability can further observed in a temporal domain. Figures A.4, A.5, A.6 plot the mean trajectories of different temporal variables sets for each outcome sub-

		Description	Units	Type	All	Cardiac	ICU	Death	Healthy
	n	Patient Count		Integer	4266	48	76	441	3701
	N	Trajectory Count		Integer	110,916	1,248	1,976	11,466	96,226
	O	Observation Count		Integer	1,286,119	14,095	20,227	184, 504	1,067,293
Vital signs									
	HR	Heart-rate	beats/minute (bpm)		82.50 (72 - 94)	88 (77 - 100)	88.50 (77 - 99)	89 (77 - 100)	81.67 (72 - 92)
	RR	Respiratory-Rate	breaths/minute (Bpm)		18 (16 - 19)	18.29 (17.50 - 20)	18 (16 - 19)	19 (18 - 21)	18 (16 - 18)
	SBP	Systolic Blood Pressure	mmHg		126 (112 - 141)	122 (108 - 138)	119 (104 - 137)	123 (108 - 140)	127 (113 - 142)
	DBP	Diastolic Blood Pressure	mmHg	Continuous	67 (60 - 76)	65 (57 - 73)	64 (57 - 72.33)	66 (58 - 75)	67 (60 - 76)
	SPO2	Estimated Oxygen Saturation	%		95 (94 - 97)	95 (93 - 97)	95 (94 - 97)	94 (91 - 96)	95 (94 - 97)
	FIO2	Fraction of Inspired Oxygen concentration	%		21 (21.00 - 28.67)	21 (21 - 31)	21 (21 - 41)	28 (21 - 43)	21 (21 - 21)
	TEMP	Temperature	°C		36.40 (36.05 - 36.80)	36.25 (36.00 - 36.60)	36.40 (36.00 - 36.83)	36.20 (36. - 36.65)	36.40 (36.10 - 36.80)
	AVPU	Alert, Verbal, Pain, Unresponsive Scale		Categorical (1-4)	1 (1 - 1)	1 (1 - 1)	1 (1 - 1)	1 (1 - 1)	1 (1 - 1)
Static									
	age	Patient age	year	Integer	72 (62 - 81)	76 (69 - 82)	69 (61 - 80)	81 (74 - 88)	71 (61 - 74)
	gender	Male patients			2123 (49.77 %)	28 (58.33 %)	38 (50.00 %)	247 (56.01 %)	1810 (48.01 %)
	Elective	Elective Admissions		Binary	1139 (26.70 %)	2 (4.17 %)	8 (10.53 %)	3 (0.68 %)	1126 (30.42 %)
	Surgical	Surgical admissions			1130 (26.49 %)	6 (12.50 %)	22 (28.95 %)	48 (10.88 %)	1054 (28.48 %)
Serum									
	HGB	Haemaglobin	g/L		11.20 (9.70 - 12.80)	11.00 (9.35 - 12.40)	10.35 (9.00 - 12.03)	10.80 (9.30 - 12.60)	11.30 (9.80 - 12.90)
	WBC	White Blood Cell count (blood)	x10^9/L		10.03 (7.63 - 13.23)	10.90 (8.41 - 15.36)	10.41 (6.79 - 14.04)	11.38 (8.34 - 15.82)	9.80 (7.54 - 12.82)
	EOS	EOSinophil count (blood)	x10^9/L		0.10 (0.02 - 0.22)	0.06 (0.01 - 0.20)	0.04 (0.00 - 0.15)	0.03 (0.00 - 0.11)	0.11 (0.03 - 0.24)
	BAS	BASophil count (blood)	x10^9/L	Continuous	0.04 (0.02 - 0.05)	0.03 (0.02 - 0.06)	0.03 (0.02 - 0.05)	0.03 (0.02 - 0.05)	0.04 (0.02 - 0.06)
	EBR	Eosinophil-Basophil Ratio			2.60 (0.75 - 5.50)	1.50 (0.50 - 4.88)	1.43 (0.17 - 4.50)	1.00 (0.00 - 3.33)	3.00 (1.00 - 5.86)
	NEU	NEUtrophil count (blood)	x10^9/L		7.60 (5.34 - 10.70)	8.88 (6.25 - 12.60)	8.14 (4.75 - 11.25)	9.31 (6.46 - 13.52)	7.31 (5.19 - 10.19)
	LYM	LYMphocyte count (blood)	x10^9/L		1.16 (0.77 - 1.68)	1.00 (0.63 - 1.58)	1.00 (0.55 - 1.65)	0.84 (0.55 - 1.24)	1.24 (0.85 - 1.75)
	NLR	Neutrophil-Lymphocyte Ratio			6.36 (3.76 - 11.67)	9.40 (4.88 - 15.98)	8.41 (4.33 - 15.55)	10.84 (6.20 - 19.59)	5.74 (3.51 - 10.11)
Haematological									
	ALB	ALBumin level (plasma)	g/L		26.00 (22.00 - 30.00)	23.00 (20.00 - 28.00)	23.00 (19.00 - 27.50)	23.00 (19.00 - 28.00)	27.00 (23.00 - 31.00)
	CR	Creatinine level (plasma)	umol/L		77.00 (58.00 - 109.00)	107.00 (78.75 - 154.25)	78.00 (52.00 - 125.75)	93.00 (63.00 - 145.00)	74.00 (58.00 - 102.00)
	CRP	C-Reactive Protein level (plasma)	mg/L	Continuous	63.43 (21.30 - 137.58)	51.50 (25.80 - 112.85)	113.20 (32.45 - 226.05)	85.20 (36.20 - 156.35)	58.00 (18.68 - 131.33)
	POT	POTassium level (plasma)	mmol/L		4.00 (3.60 - 4.40)	4.20 (3.80 - 4.80)	4.10 (3.70 - 4.50)	4.10 (3.70 - 4.80)	4.00 (3.60 - 4.30)
	SOD	SODium level (plasma)	mmol/L		137.00 (134.00 - 140.00)	136.00 (132.00 - 139.00)	136.00 (133.00 - 139.00)	138.00 (134.00 - 142.00)	137.00 (135.00 - 140.00)
	UR	URea concentration levels	mL		6.40 (4.40 - 10.40)	10.50 (6.45 - 18.15)	7.05 (4.40 - 11.67)	10.40 (6.40 - 17.50)	5.90 (4.20 - 9.00)

A.2: Descriptive statistics and information of all input data features. Variables are displayed with type, description, units and average statistics. We separate all features according to medical literature, including vital-sign, static, serum and haematological variables, and we also display statistics per outcome sub-cohort, defined as a cohort with those patients assigned to a given outcome.

	No Event	Death	ICU	Cardiac
N	3701	441	76	48
Age (IQR)	71 (61 - 80)	81 (74 - 88)	69 (61 - 74)	76 (69 - 82)
Gender, M	1810 (48.9%)	247 (56.0%)	38 (50.0%)	28 (58.33%)
CCI (IQR)	4 (3 - 13)	14 (4 - 21)	7 (4 - 17)	15 (4 - 23)
Elective	1126 (30.4%)	3 (0.7%)	8 (10.5%)	2 (4.2%)
Surgical	1054 (28.5%)	48 (10.1%)	22 (29.0%)	6 (12.5%)

A.3: Descriptive demographic variable information for each outcome sub-cohort.

cohort, respectively, according to vital signs, haematological and serum features. Mean is calculated based on the time to outcome, and missing observations are disregarded and ignored.

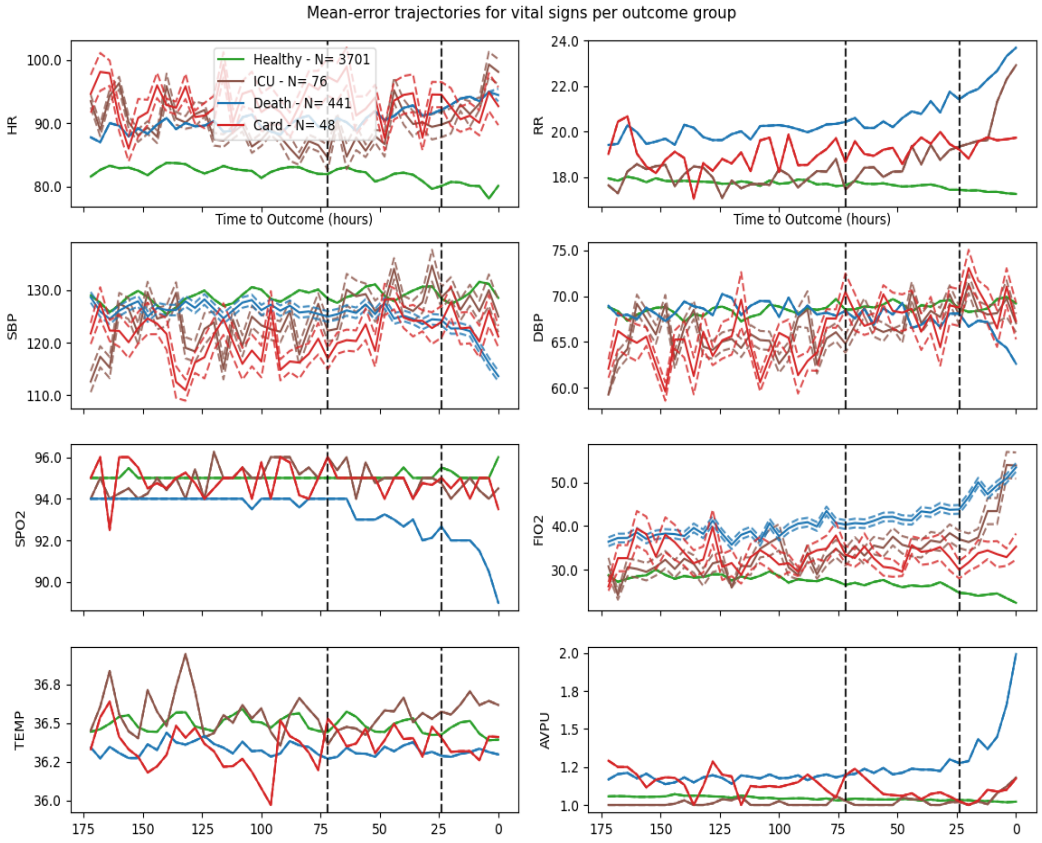

A.4: Plot of mean vital-sign trajectories (median with respect to SpO_2) in solid line as given by the 4 outcome groups: admissions with a) Cardiac-, b) Death- , or c) ICU-, and d) No-events. The respective standard errors are represented by the dashed lines. We visualised trajectories from up to 7 days prior to an outcome event or discharge - the black lines represent the time window (72 - 24 hours prior to an event or discharge) considered for input to all models.

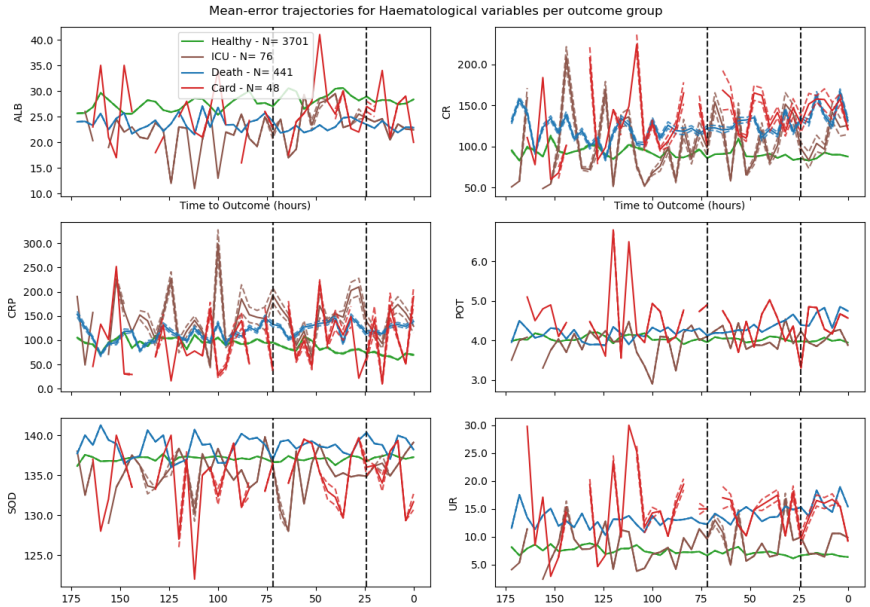

A.5: Plot of mean haematological trajectories in solid line as given by the 4 outcome groups: admissions with a) Cardiac-, b) Death- , or c) ICU-, and d) No-events. The respective standard errors are represented by the dashed lines. We visualised trajectories from up to 7 days prior to an outcome event or discharge - the black lines represent the time window (72 - 24 hours prior to an event or discharge) considered for input to all models.

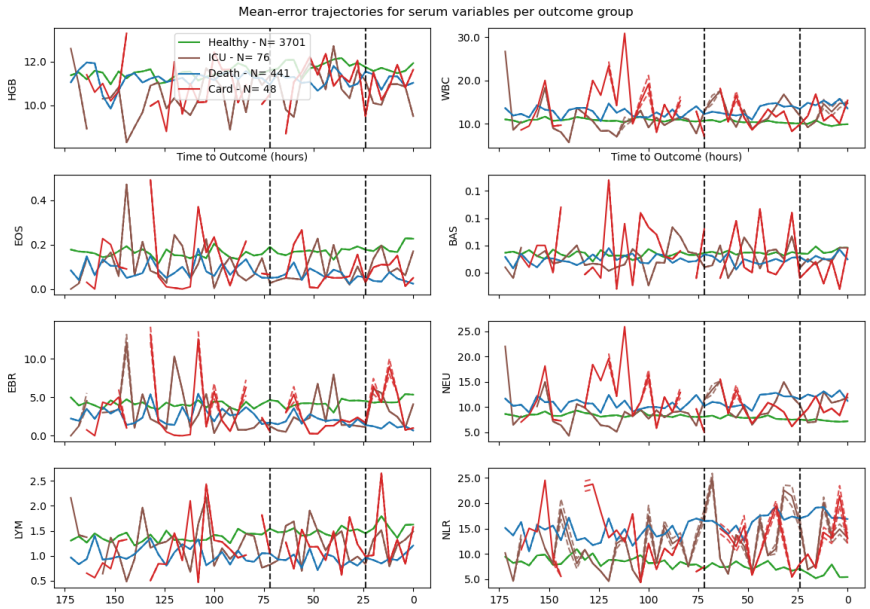

A.6: Plot of mean serum trajectories in solid line as given by the 4 outcome groups: admissions with a) Cardiac-, b) Death- , or c) ICU-, and d) No-events. The respective standard errors are represented by the dashed lines. We visualised trajectories from up to 7 days prior to an outcome event or discharge - the black lines represent the time window (72 - 24 hours prior to an event or discharge) considered for input to all models.

A list indicating the grid-search range of hyper-parameters considered in our experiments are indicated in Table A.7. For simplicity, we define $\mathbb{P} := \{0.001, 0.01, 0.1, 1, 10\}$, $\mathbb{L} := \{32, 64, 128, 256\}$ and $\mathbb{K} := \{3, ..., 20\}$. In bold, top-performing hyper-parameters according to target metrics defined in Section 4 are highlighted.

Parameter	TSKM	SOM-VAE	AC-TPC	CAMELOT	SVM	XGB
seeds	$\{1001, 1012, 1134, 2475, 6138, 7415, 1663, 7205, 9253, 1782\}$					
α	-	$\mathbb{P}$ (**0.1**)	$\mathbb{P}$ (**0.01**)	$\mathbb{P}$ (**0.01**)	-	-
β	-	$\mathbb{P}$ (**0.1**)	$\mathbb{P}$ (**0.01**)	$\mathbb{P}$ (**0.001**)	-	-
γ	-	-	-	-	-	$\{0.1, \mathbf{0.2}, 0.6\}$
latent dim	-	$\mathbb{L}$ (**64**)	$\mathbb{L}$ (**128**)	$\mathbb{L}$ (**128**)	-	-
SOM dim	-	$\mathbb{L}^2$(**4**, **4**)	-	-	-	-
K	$\mathbb{K}$(**7**)	-	$\mathbb{K}$ (**6**)	$\mathbb{K}$ (**6**)	-	-
kernel	{**'DTW'**, 'euclidean'}	-	-	-	{**'polynomial'**, 'rbf'}	-
C	-	-	-	-	$\mathbb{P}$ (**10**)	-
n-estimators	-	-	-	-	-	$\{100, \mathbf{200}, 300\}$
depth	-	-	-	-	-	$\{1, 3, \mathbf{5}, 10\}$
min-child-weight	-	-	-	-	-	$\{1, 2, 3, \mathbf{5}\}$

A.7: Parameter range used for Grid-search hyper-parameter optimisation. For each model, the list of parameter values tested is indicated. In **bold**, the optimum set of hyper-parameters is indicated for each model.

In Figures A.8 and Figures A.9 we display cluster phenotypes for some of our experiments with benchmark clustering models SOM-VAE and TSKM, respectively. Both models do not associate clusters with phenotypes - we estimate the cluster phenotype as the *empirical* outcome distribution for the patient cohort assigned to the corresponding cluster.

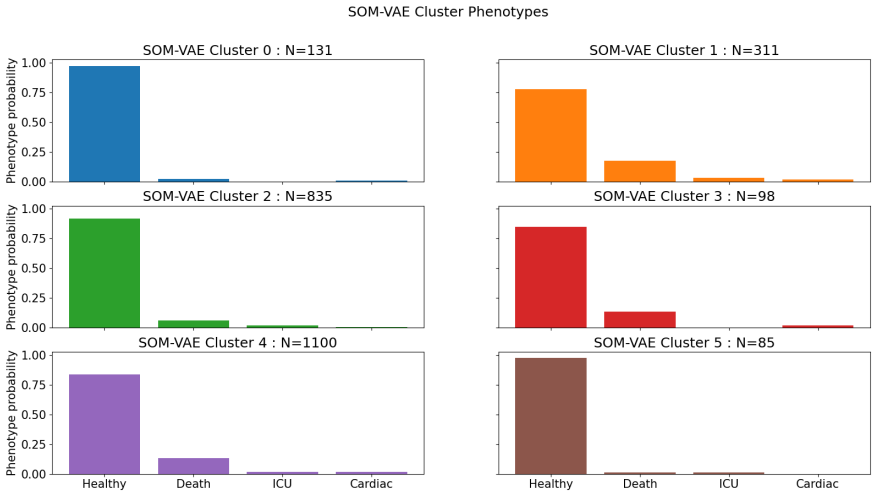

A.8: Bar plots of learnt cluster phenotypes for SOM-VAE with optimal hyper-parameters. Each plot represents a cluster - its phenotype is the corresponding empirical outcome distribution in its cluster-assigned patient cohort.

We note that clusters learnt by both clustering benchmark models have identical phenotypes, which provides no useful clinical interpretability to the cluster-defined populations, as well as likely not assisting models to learn relevant cluster representations.

A complete description of the number of patient admissions with a given outcome per learnt cluster in our proposed model can be seen in Table A.10.

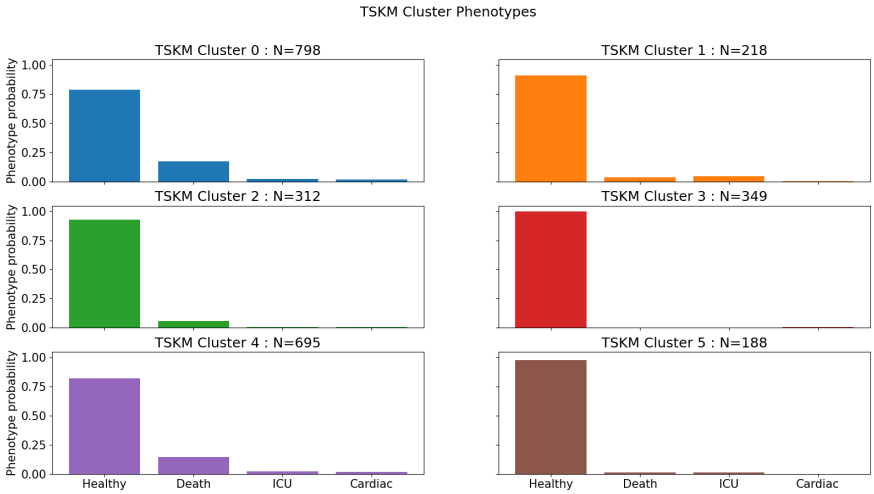

A.9: Bar plots of learnt cluster phenotypes for TSKM with $K = 6$. Each plot represents a cluster - its phenotype is the corresponding empirical outcome distribution in its cluster-assigned patient cohort.

Outcome	Healthy	Death	ICU	Cardiac
Cluster 0	33	43	1	2
Cluster 1	42	20	7	1
Cluster 2	645	2	3	1
Cluster 3	216	129	15	12
Cluster 4	339	12	10	1
Cluster 5	847	59	9	12

A.10: Table with empirical number of outcome admissions observed for each cluster learnt by the proposed model.