UNIVERSITY OF CAMBRIDGE

Representation Learning for Patients in the Intensive Care Unit

Emma Rocheteau

Department of Computer Science and Technology, University of Cambridge

Thesis Presentation

6th March, 2023

Chapter 3: Temporal Pointwise Convolution

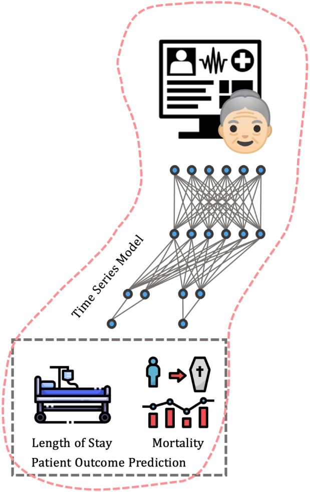

Chapter 3: Temporal Pointwise Convolution

Data: Electronic Health Records in Intensive Care

eICU

- 200,859 ICU stays
- Admitted between 2014 and 2015
- 208 different hospitals across the US

MIMIC-IV

- 69,619 ICU stays
- Admitted between 2008 and 2019
- Beth Israel Deaconess Medical Center in Boston

Data: Electronic Health Records in Intensive Care

Both datasets contain:

- Vital Signs e.g. heart rate
- Lab Results e.g. blood glucose
- Demographics e.g. age
- Diagnoses
- Medications

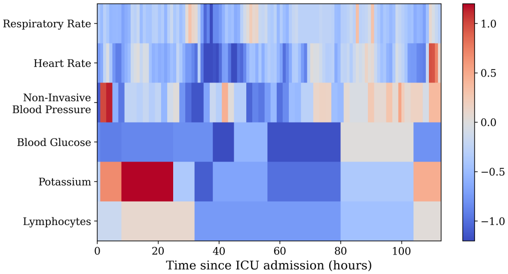

What do we want the model to extract?

- Temporal trends
- Inter-feature relationships

Example

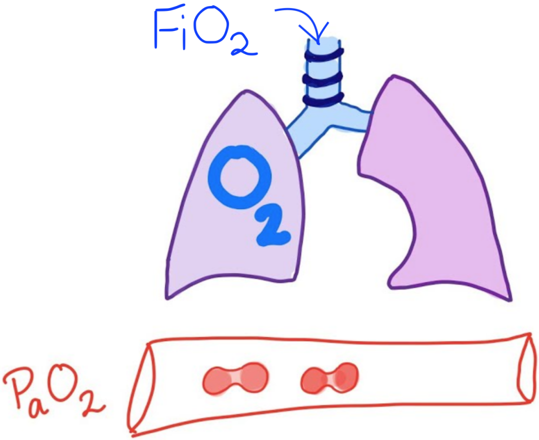

Temporal Convolution

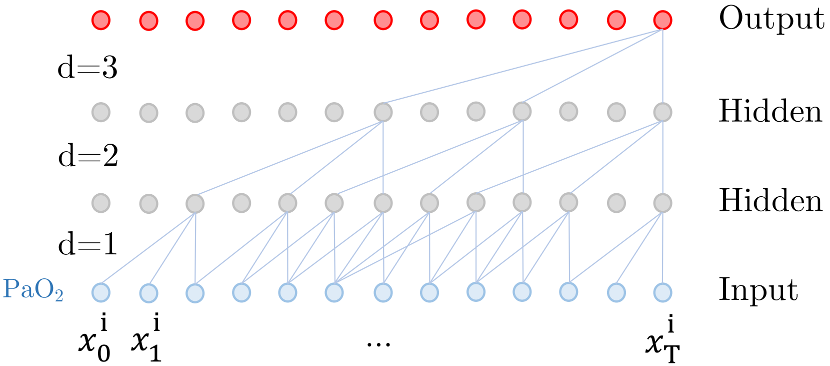

Temporal Receptive Fields

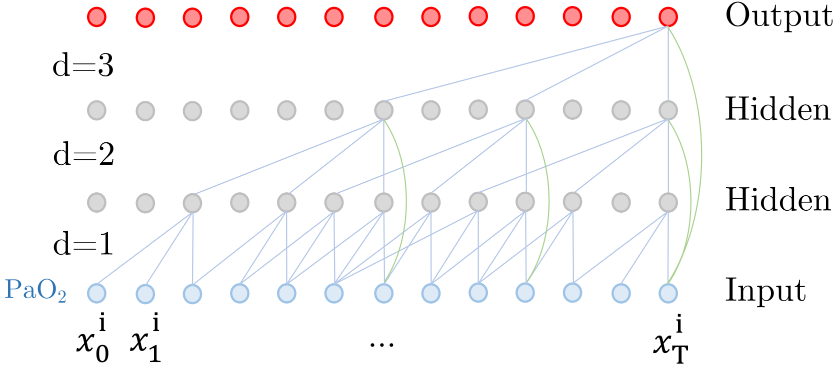

Pointwise Convolution

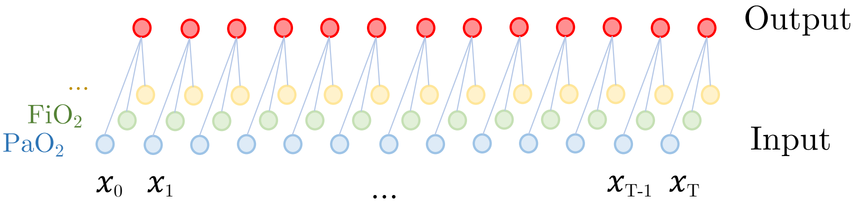

Model (one TPC layer)

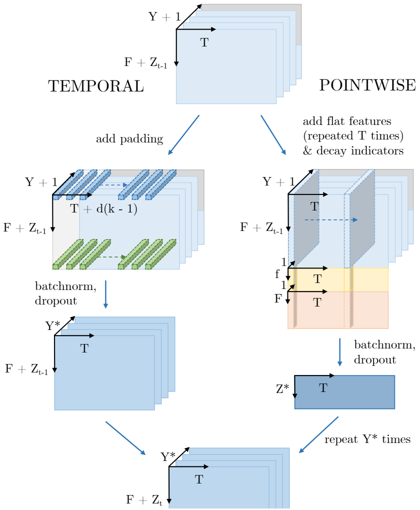

eICU LoS Results

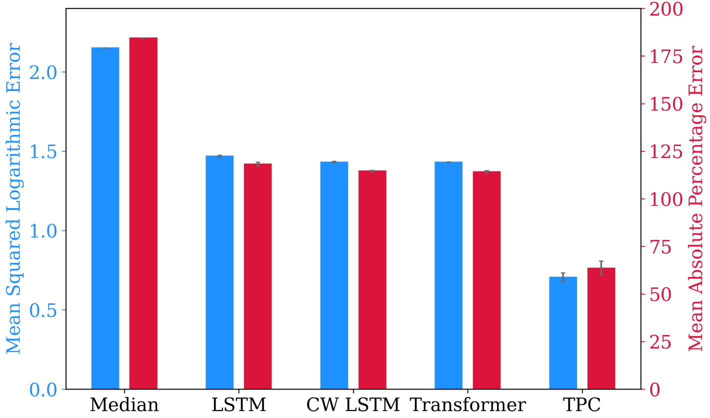

MIMIC-IV LoS Results

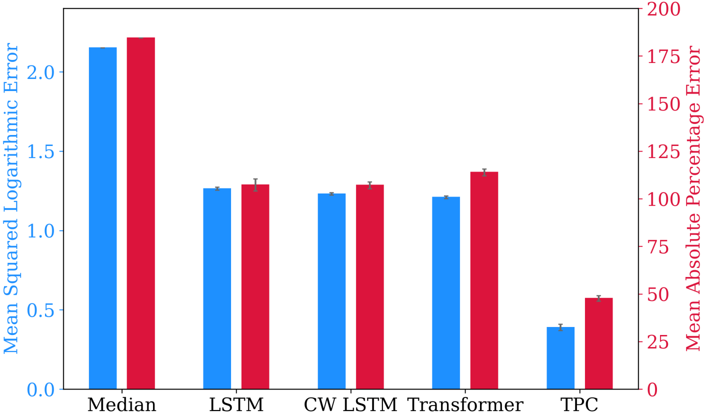

eICU Mortality Results

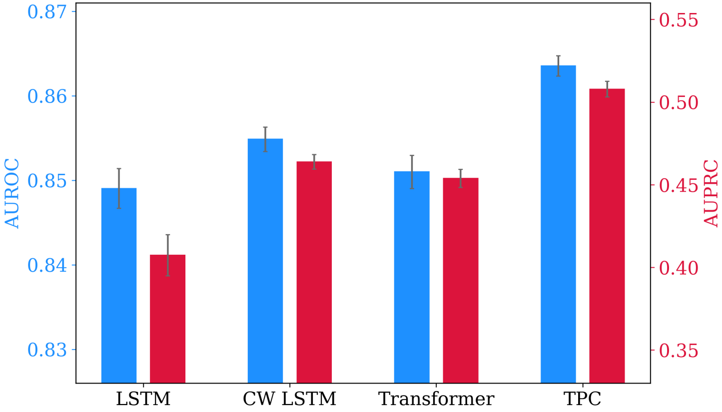

MIMIC-IV Mortality Results

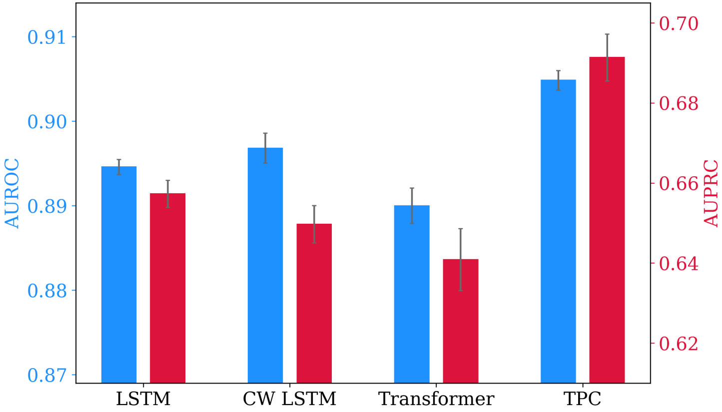

Why does TPC do well on EHR time series?

- It has been specifically designed to be able to extract trends and inter-feature relationships.
- It can choose its own temporal receptive field sizes (independently for each feature) because of the skip connections.
- Rigid convolutional filters can exploit periodicity in EHR timeseries.

Chapter 4: Graph Representation Learning

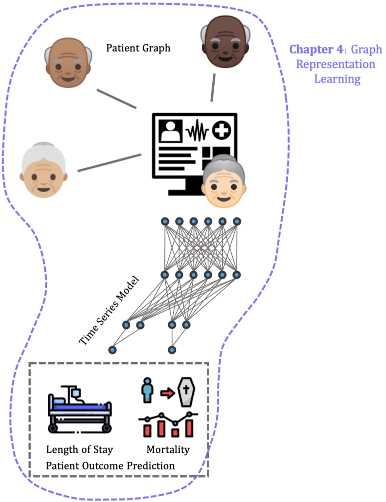

Diagnosis Information is Hard to Use

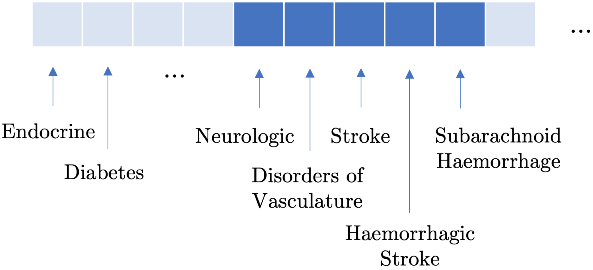

Distribution of Diagnoses in the eICU Database

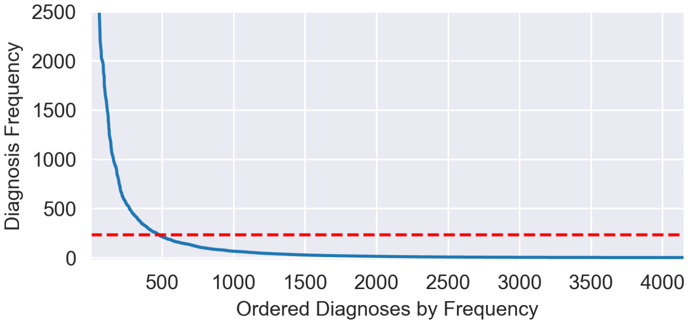

“Relatedness”: Grouping Similar Patients

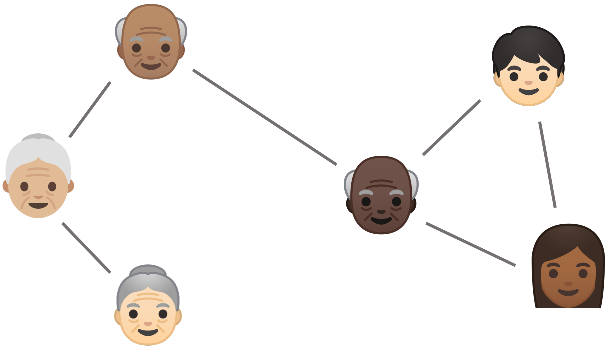

Graph Neural Networks (GNNs)

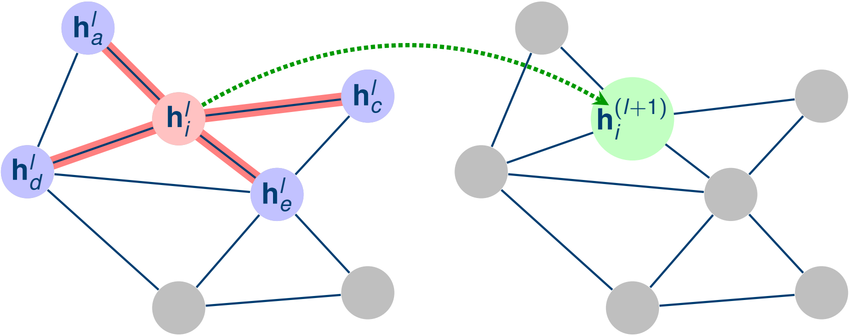

Graph Construction

The "relatedness" score between two patients *i* and *j* is given by:

$$\mathcal{M}_{ij} = a\overbrace{\sum_{\mu=1}^{m}\left(\mathcal{D}_{i\mu}\mathcal{D}_{j\mu}(d_\mu^{-1}+c)\right)}^{\text{Shared Diagnoses}} - \overbrace{\sum_{\mu=1}^{m}\left(\mathcal{D}_{i\mu}+\mathcal{D}_{j\mu}\right)}^{\text{All Diagnoses}} \tag{1}$$

where

- $\mathcal{D} \in \mathbb{R}^{N\times m}$ is a diagnosis matrix,
- N is the number of patients,
- m is the number of unique diagnoses,
- d_μ is the frequency of diagnosis μ,
- a and c are tunable constants.

Hybrid LSTM-GNN Model

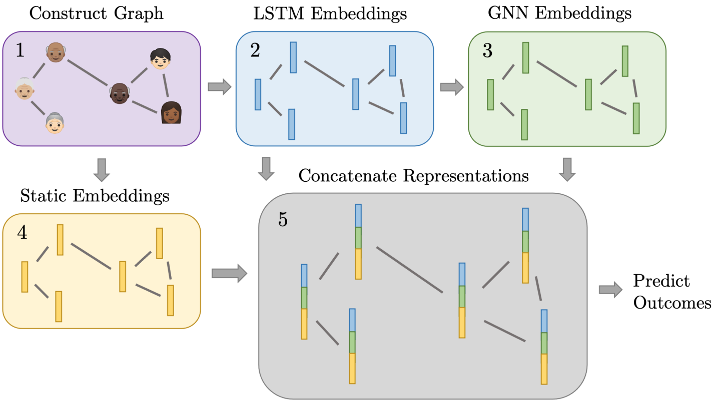

Results

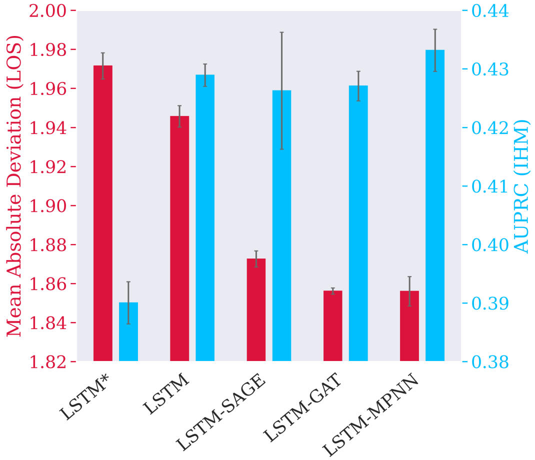

Mean Absolute Deviation (LOS)
AUPRC (IHM)
2.00
1.98
1.96
1.94
1.92
1.90
1.88
1.86
1.84
1.82
0.44
0.43
0.42
0.41
0.40
0.39
0.38
LSTM*
LSTM
LSTM-SAGE
LSTM-GAT
LSTM-MPNN

Visualisation: LSTM-GAT* attention weights

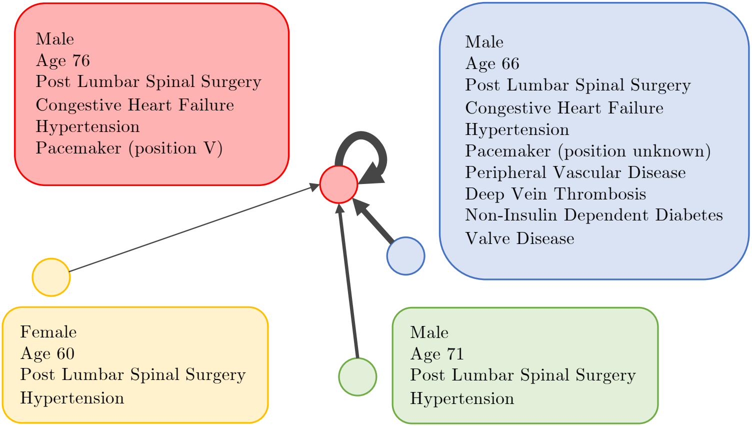

Chapter 5: Dynamic Outcomes-Based Clustering

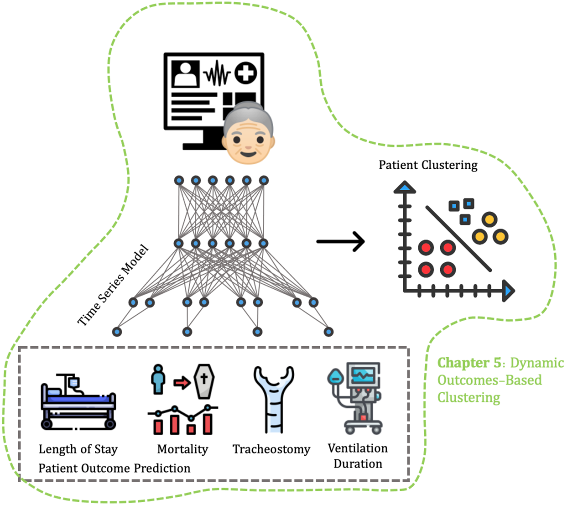

Why Cluster Patients on Mechanical Ventilation?

- Patients on mechanical ventilation are highly heterogeneous.
- Clustering would help to generate:
  - Interpretable early warning systems.
  - Further understanding of disease trajectories.
  - Early categorisation of patients for intervention.

Data: Electronic Health Records in Intensive Care

AmsterdamUMCdb

- 14,836 ventilation episodes.
- Contains:
  - Vital Signs e.g. heart rate, blood pressure
  - Lab Results e.g. blood glucose
  - Demographics e.g. age, sex, ethnicity

Methods

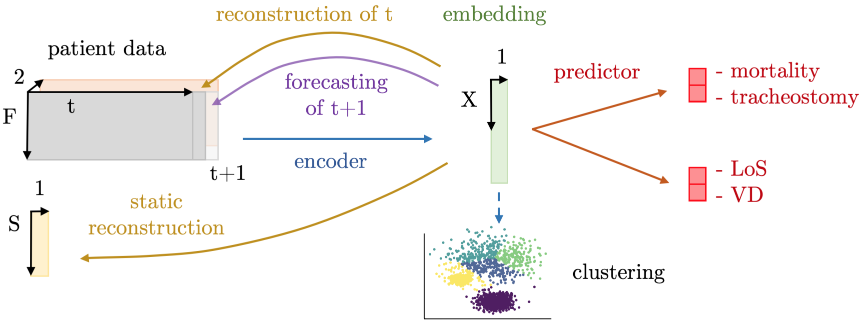

reconstruction of t
embedding
patient data
2
F
t
1
X
predictor
- mortality
- tracheostomy
forecasting
of t+1
encoder
t+1
- LoS
- VD
1
S
static
reconstruction
clustering

Outcome Task Performance: Binary

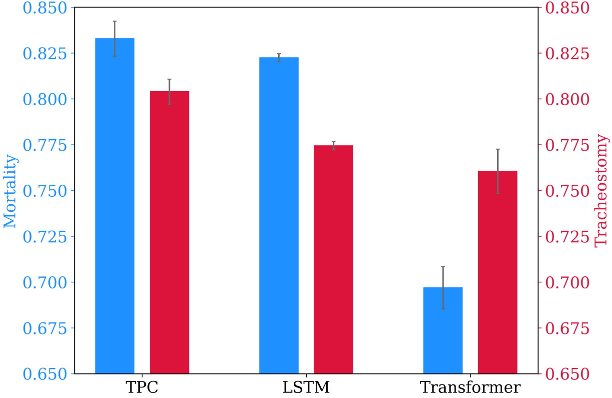

Outcome Task Performance: Duration

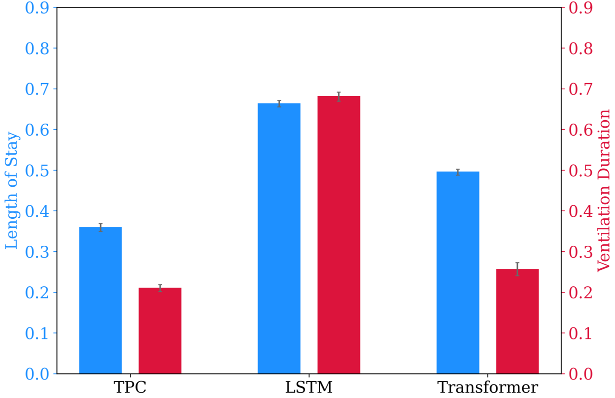

Cluster Analysis

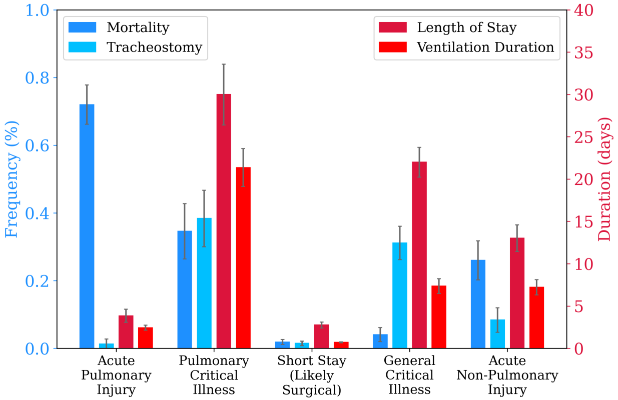

Latent Space Visualisation

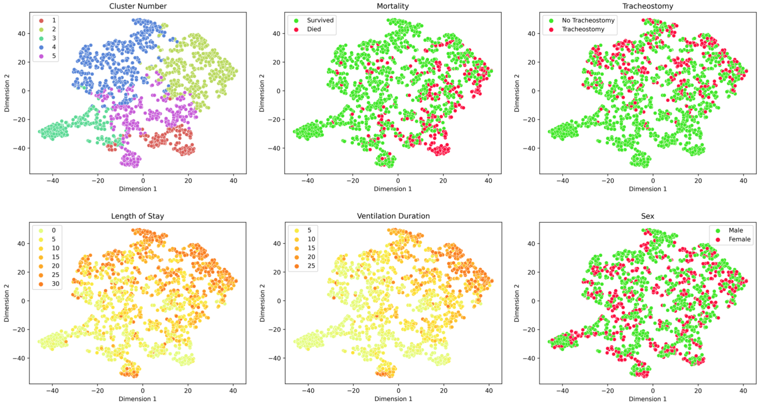

The clusters are remarkably stable over time

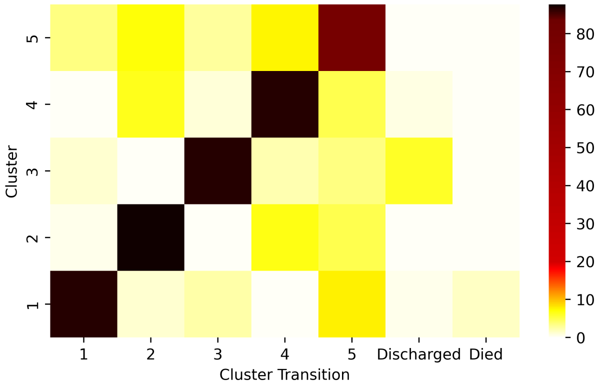

Most patients only appear in one cluster

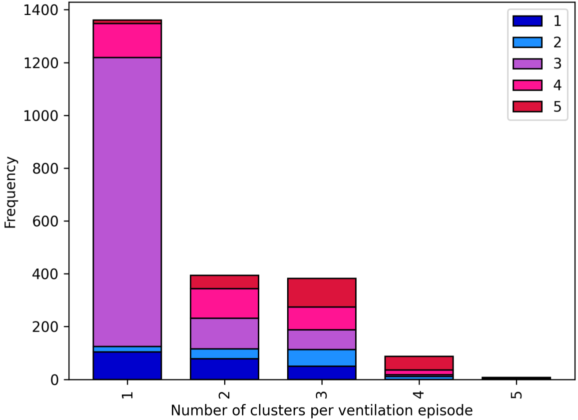

Stable categorisation happens very early

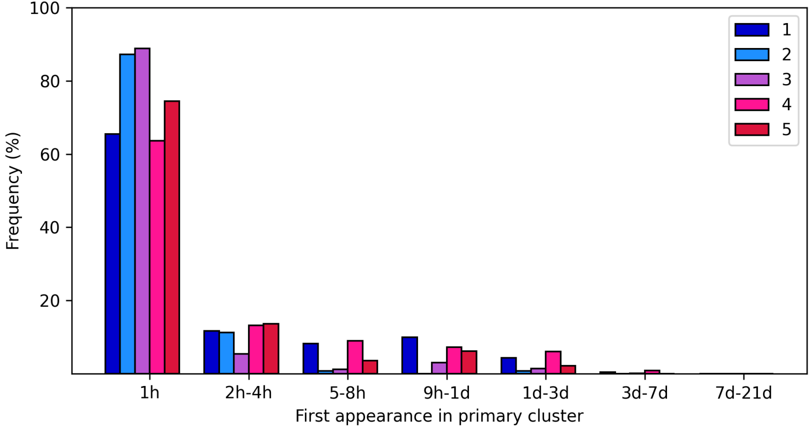

Summary

1. The TPC model outperforms alternative encoders.
2. We can generate clinically meaningful and interpretable clusters.
3. The clusters are remarkably stable across time, and membership is determined early on.
4. Stable cluster transitions do occur, and are an important avenue for future work.

Thank you! (With special mentions to...)

My funders:

The Armstrong Fund and The Frank Elmore Fund

My supervisor:

Pietro Liò

My wonderful co-authors and mentors:

Stephanie Hyland, Petar Veličković, Catherine Tong, Ioana Bica, Nicholas Lane, Rudolf Cardinal and Ari Ercole

Model Reliability

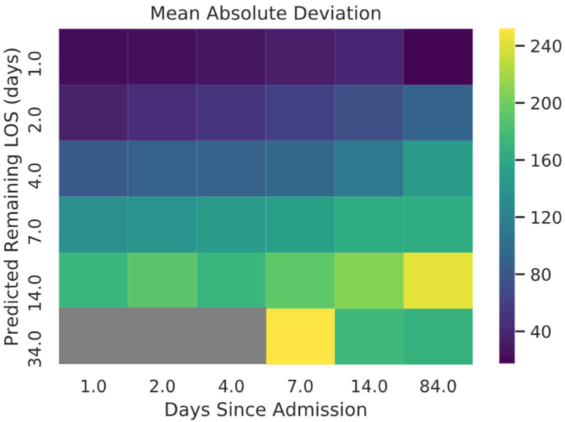

ICU Simulation Study

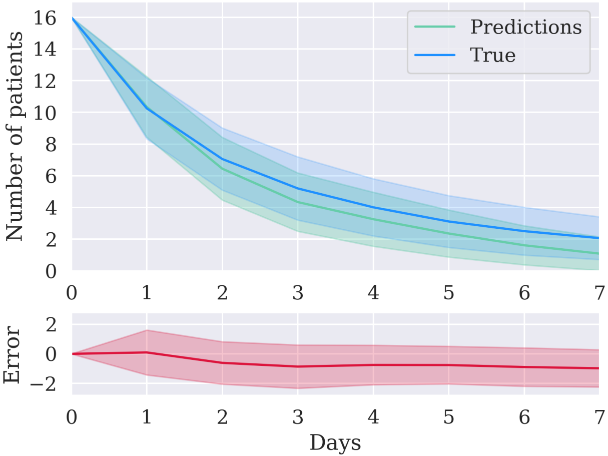

Data Type Ablation

Training Data Size

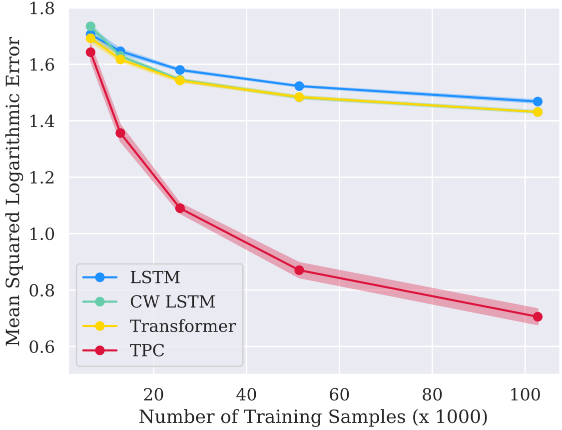

Alternative Encoders

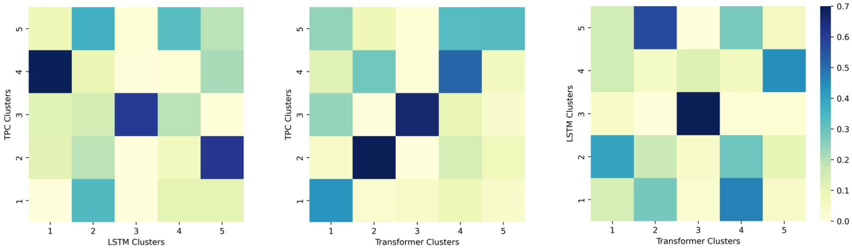

Different Initialisation Seeds

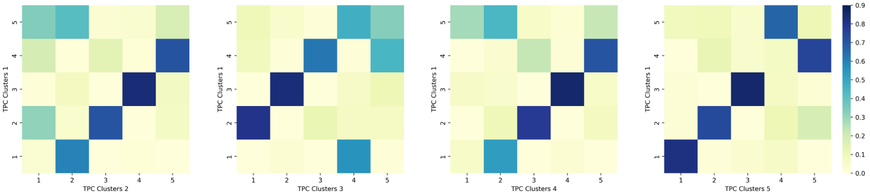

Number of Clusters

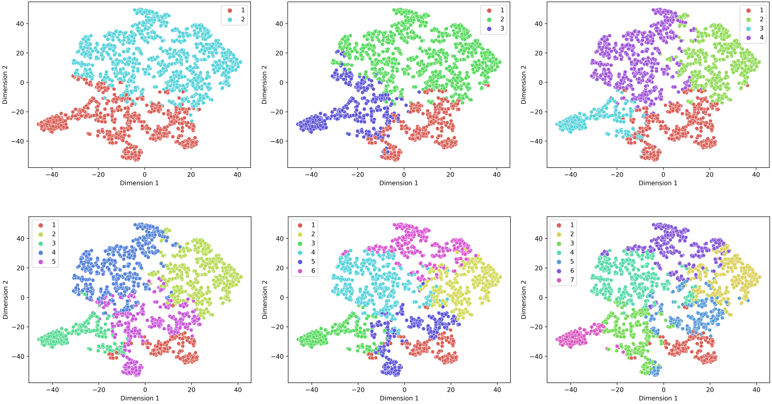

"Stable" Cluster Transitions

Transition	Count	Median Time	Mortality (%)	Tracheostomy (%)	Urgency (%)	VD	LoS
3→1	17	3	76.5	0.0	47.1	0.5	0.7
5→1	29	16	51.7	10.3	55.2	4.3	5.3
1→3	28	11	10.7	0.0	67.9	1.0	2.6
5→3	46	9	15.2	4.3	41.3	1.2	6.5
2→4	28	17	10.7	21.4	42.9	6.2	12.8
5→4	27	10	11.1	7.4	48.1	3.4	9.1
1→5	25	3	44.0	4.0	68.0	3.9	6.5
3→5	15	4	13.3	13.3	53.3	1.9	4.6
4→5	15	56	26.7	26.7	46.7	6.6	11.5