UNIVERSITY OF CAMBRIDGE

Predicting Length of Stay with Temporal Pointwise Convolutional Networks

Emma Rocheteau[1], Prof. Pietro Liò[1], and Dr Stephanie Hyland[2]

[1]Department of Computer Science and Technology, University of Cambridge, UK
[2]Microsoft Research, Cambridge, UK

ACM Conference on Health, Inference, and Learning **8-9th April, 2021**

Length of Stay

Why Length of Stay?

- It's a key determinant of hospital costs.
- Long stays increase the risk of hospital acquired infections.

Efficient bed management could mitigate costs and risk. We need to know how long the patients are going to remain in the ICU.

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

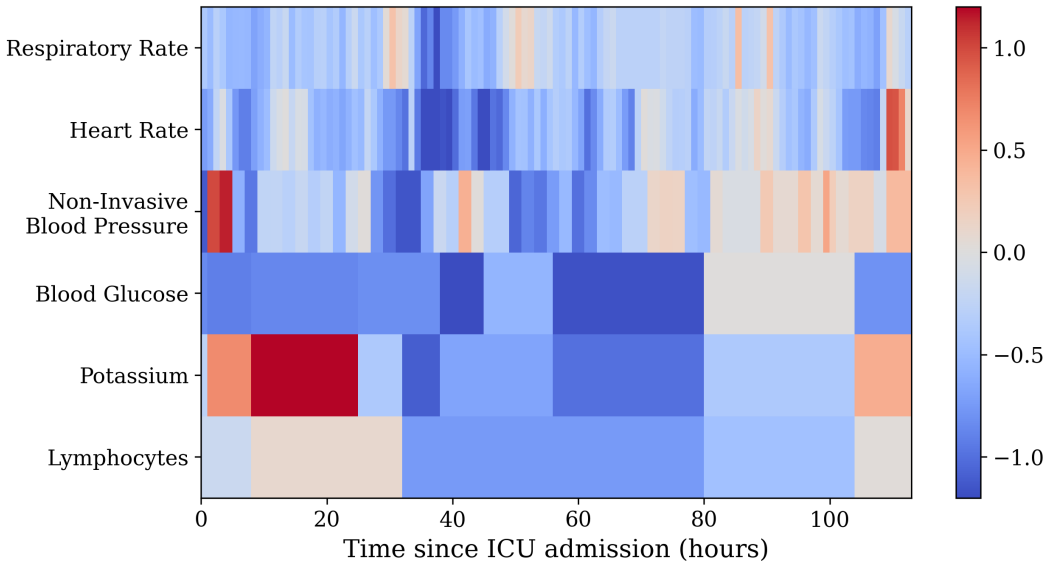

What do we want the model to extract?

- Temporal trends
- Inter-feature relationships

Example

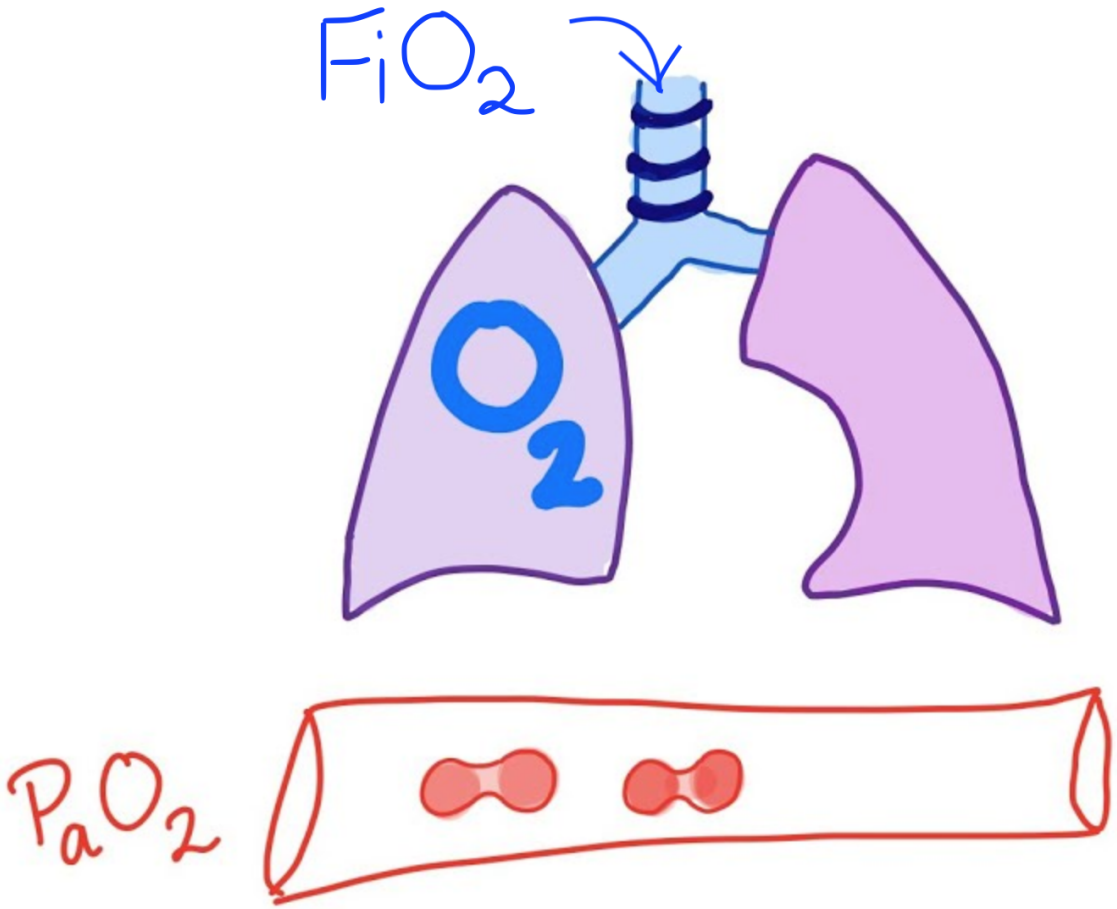

Temporal Convolution

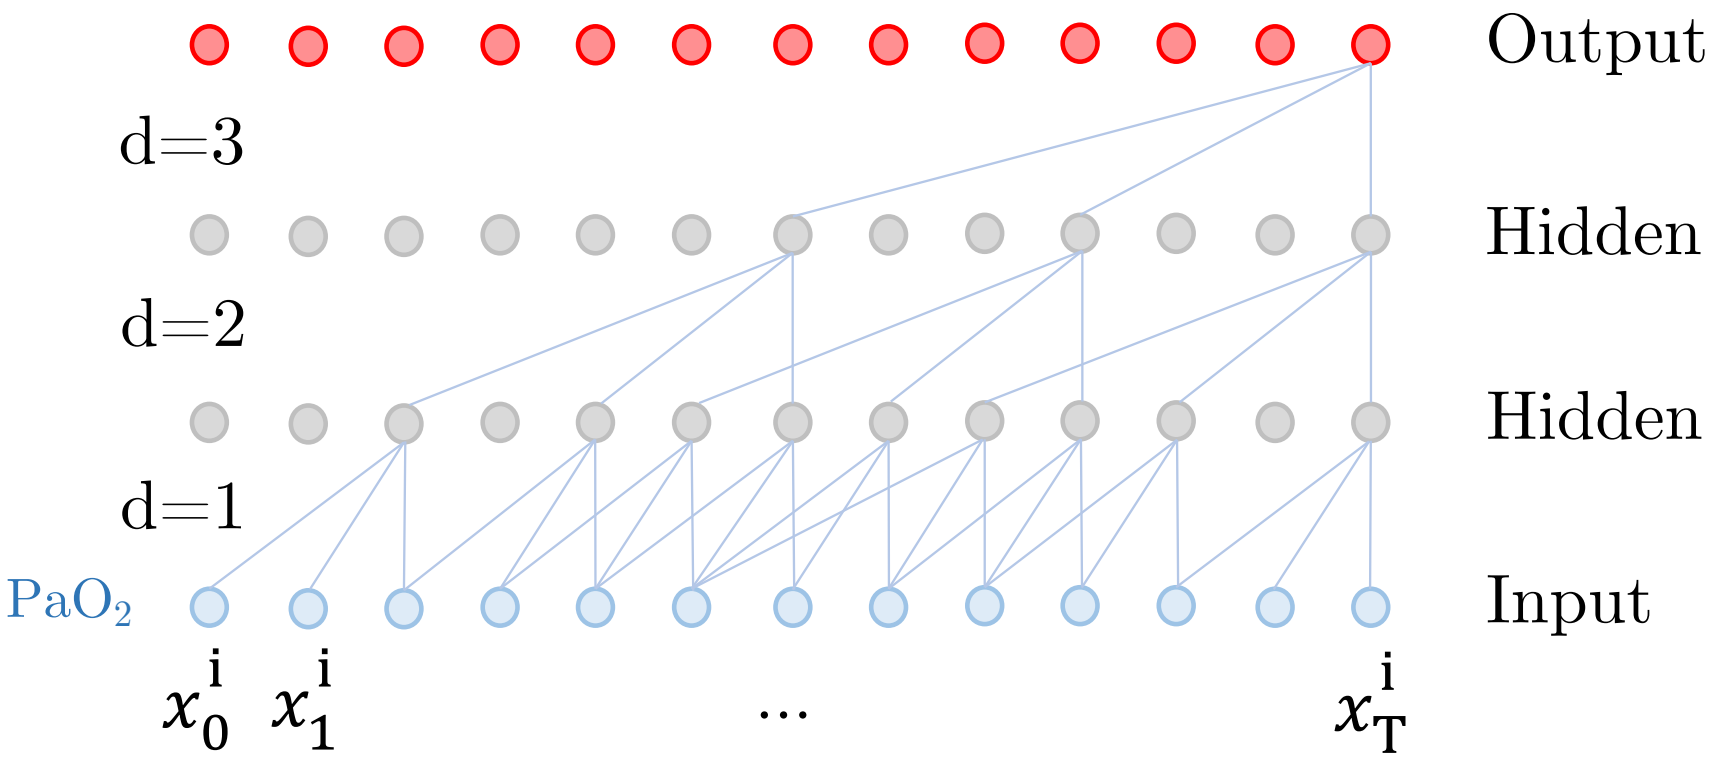

Temporal Receptive Fields

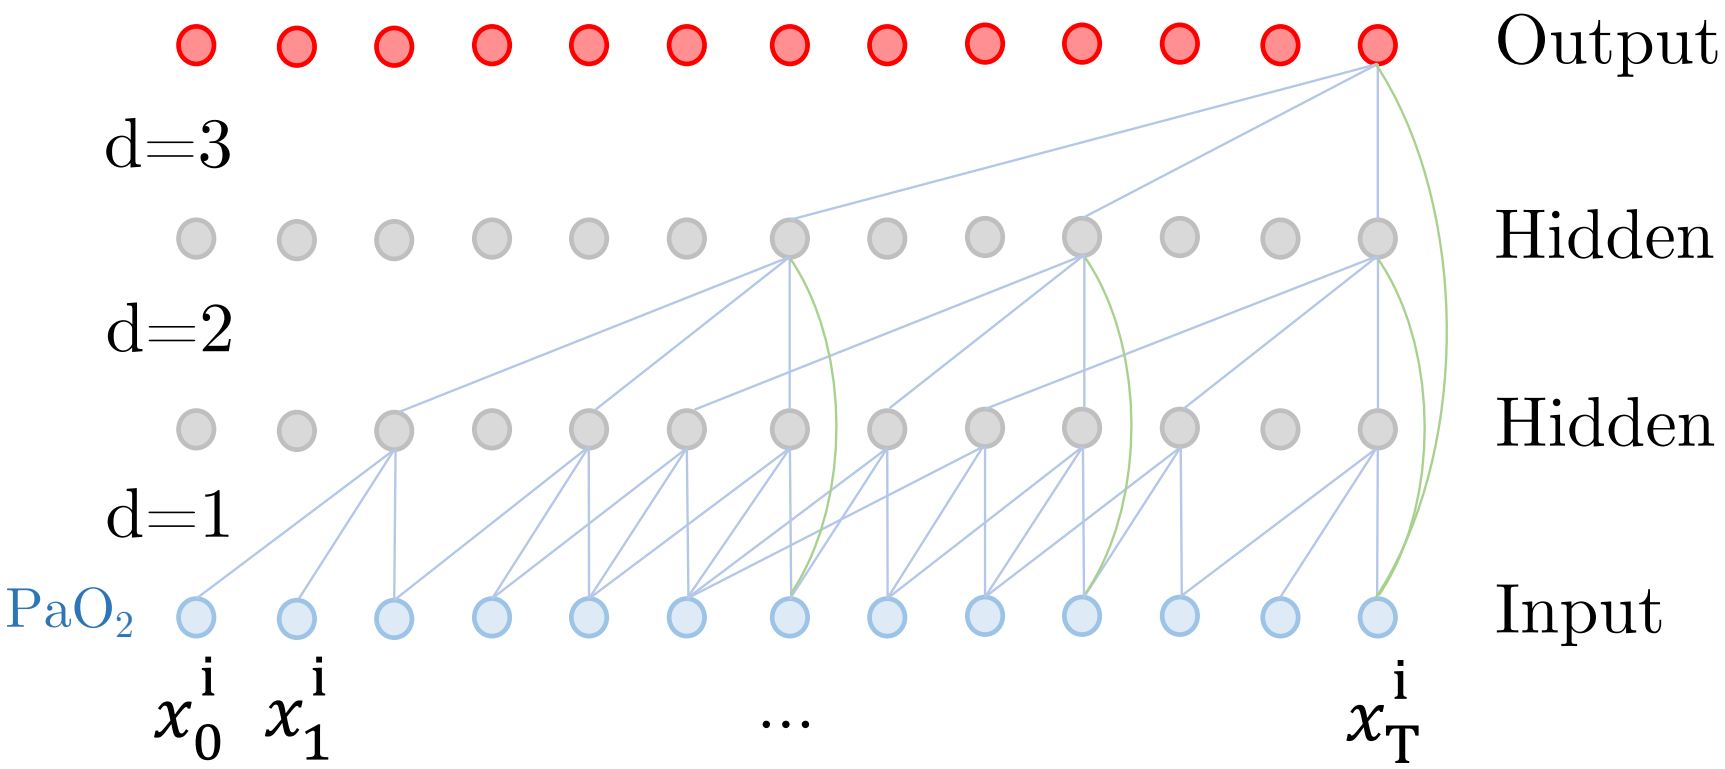

Pointwise Convolution

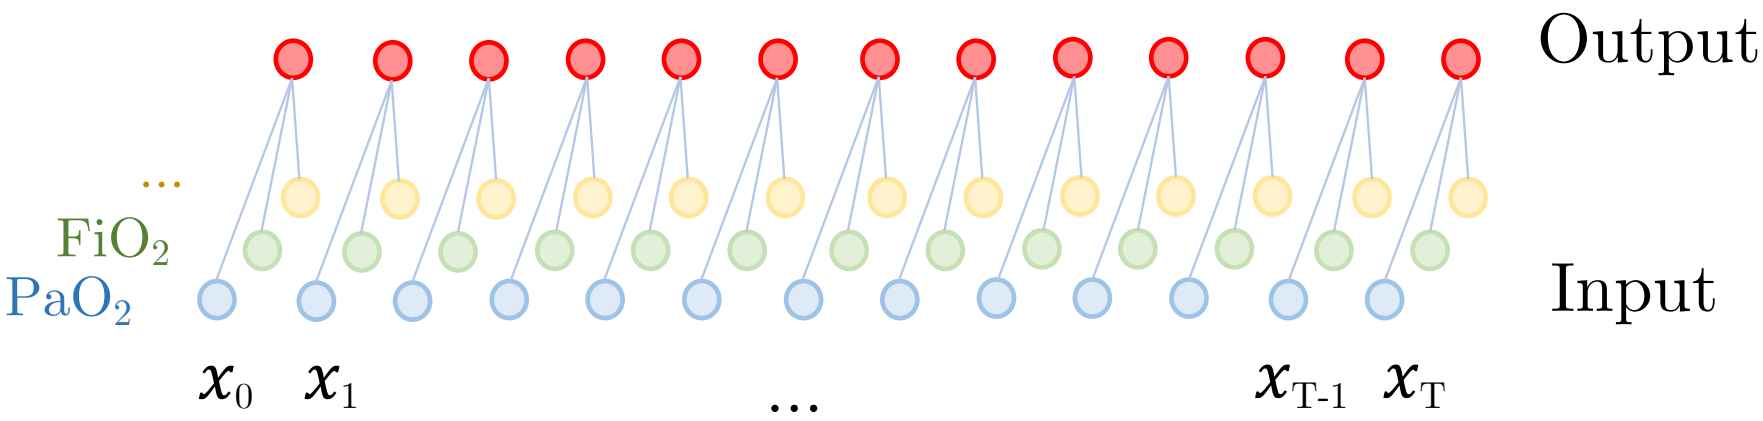

Model (one TPC layer)

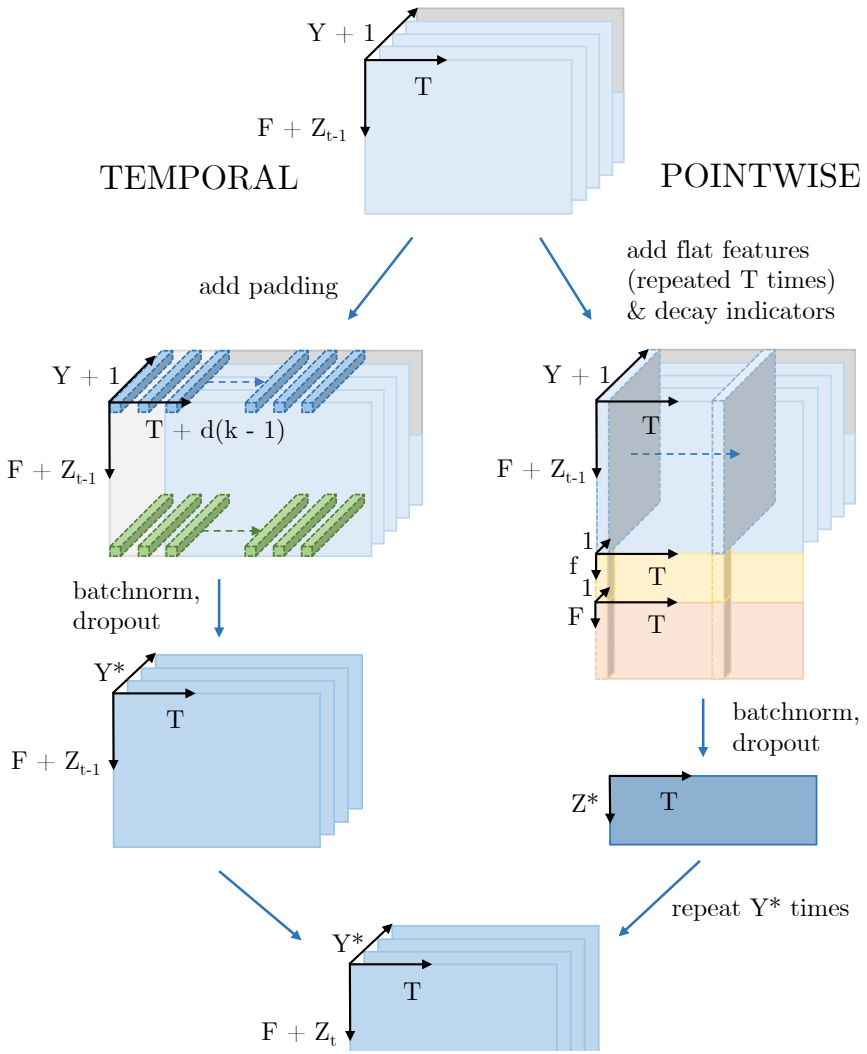

eICU LoS Results

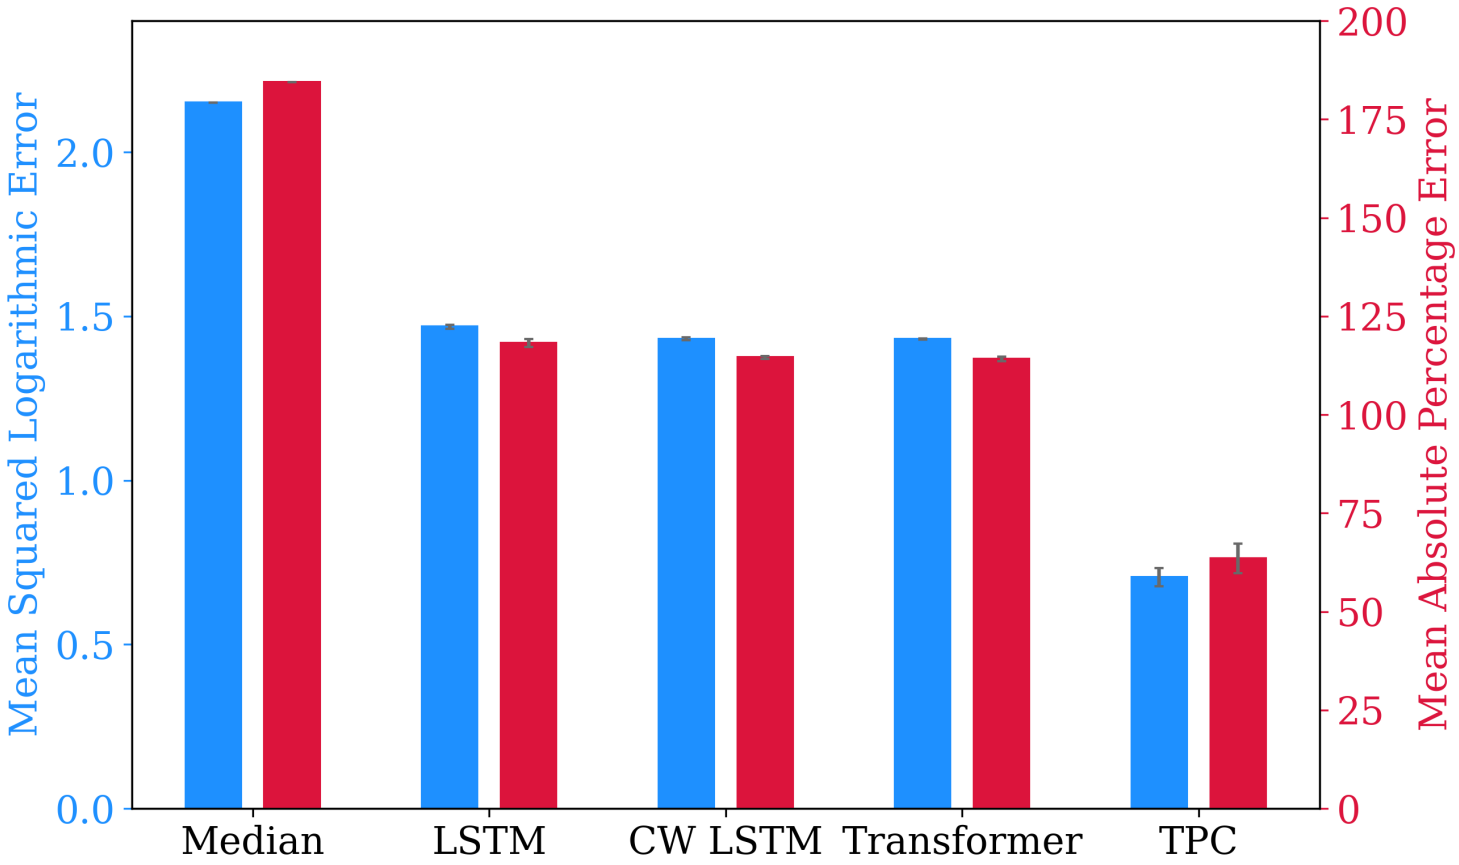

MIMIC-IV LoS Results

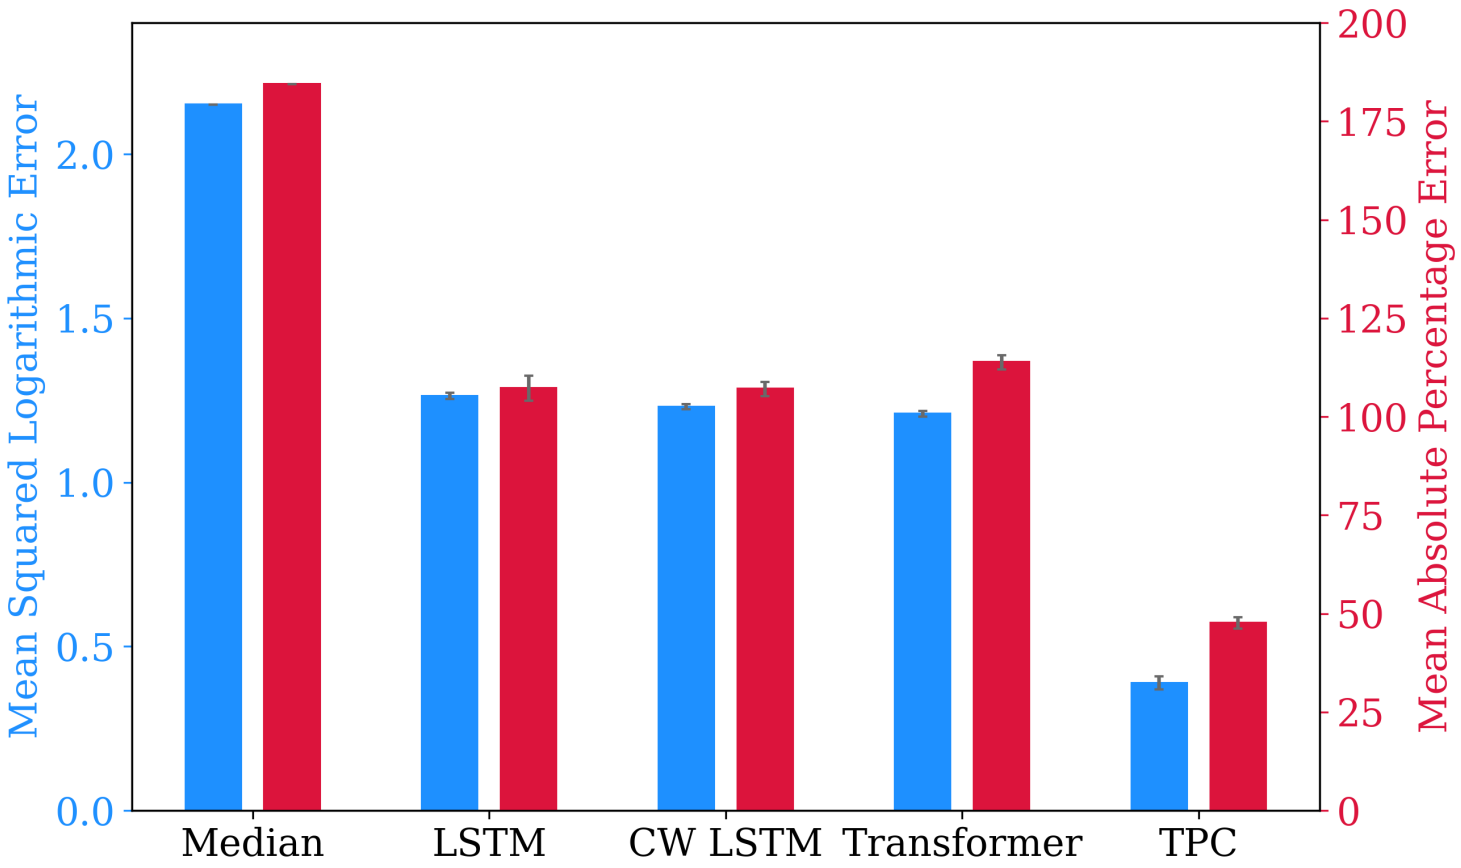

eICU Mortality Results

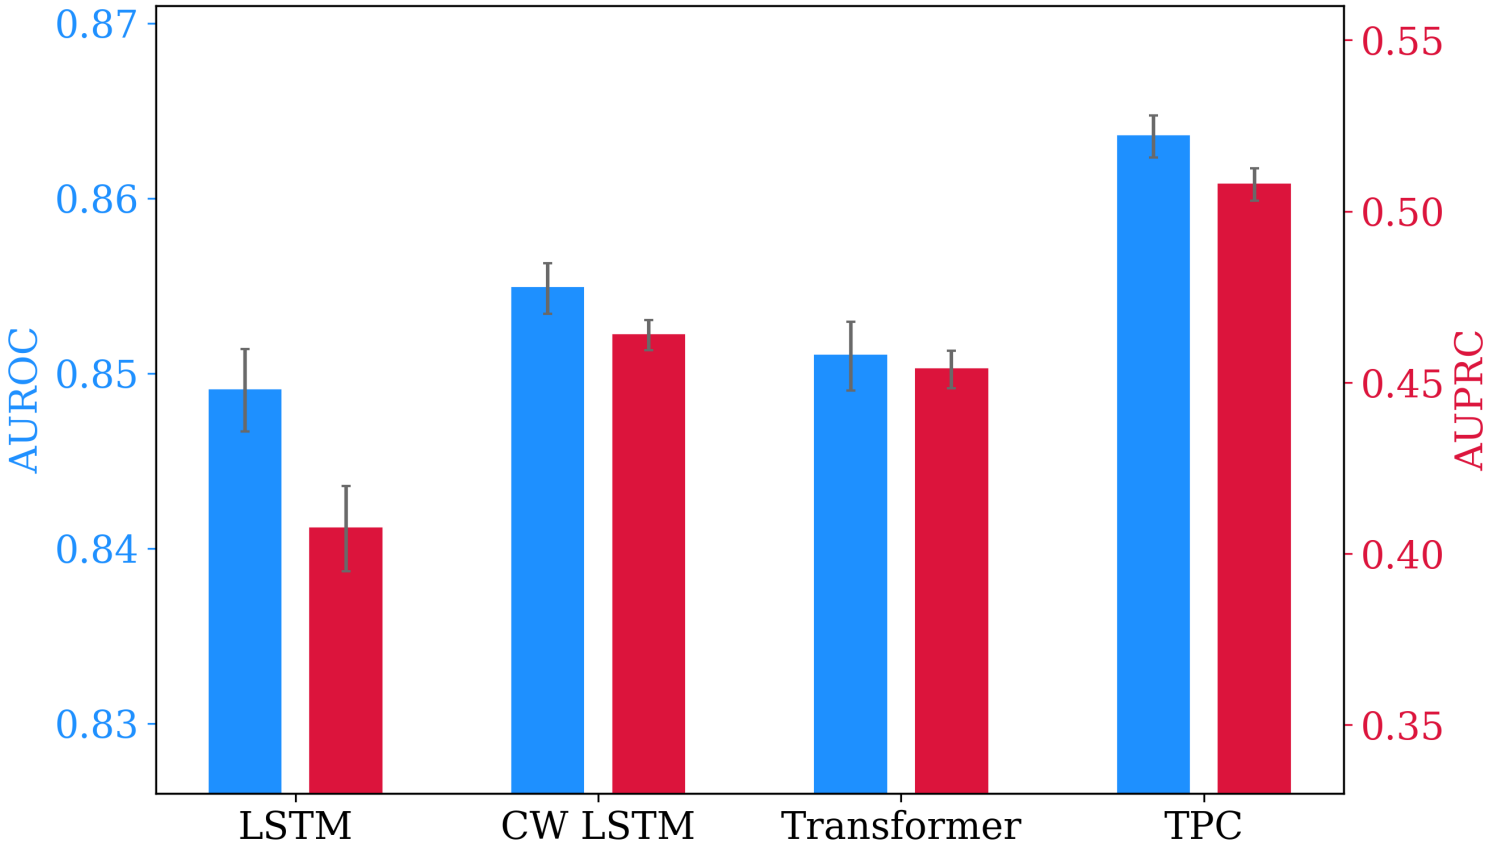

MIMIC-IV Mortality Results

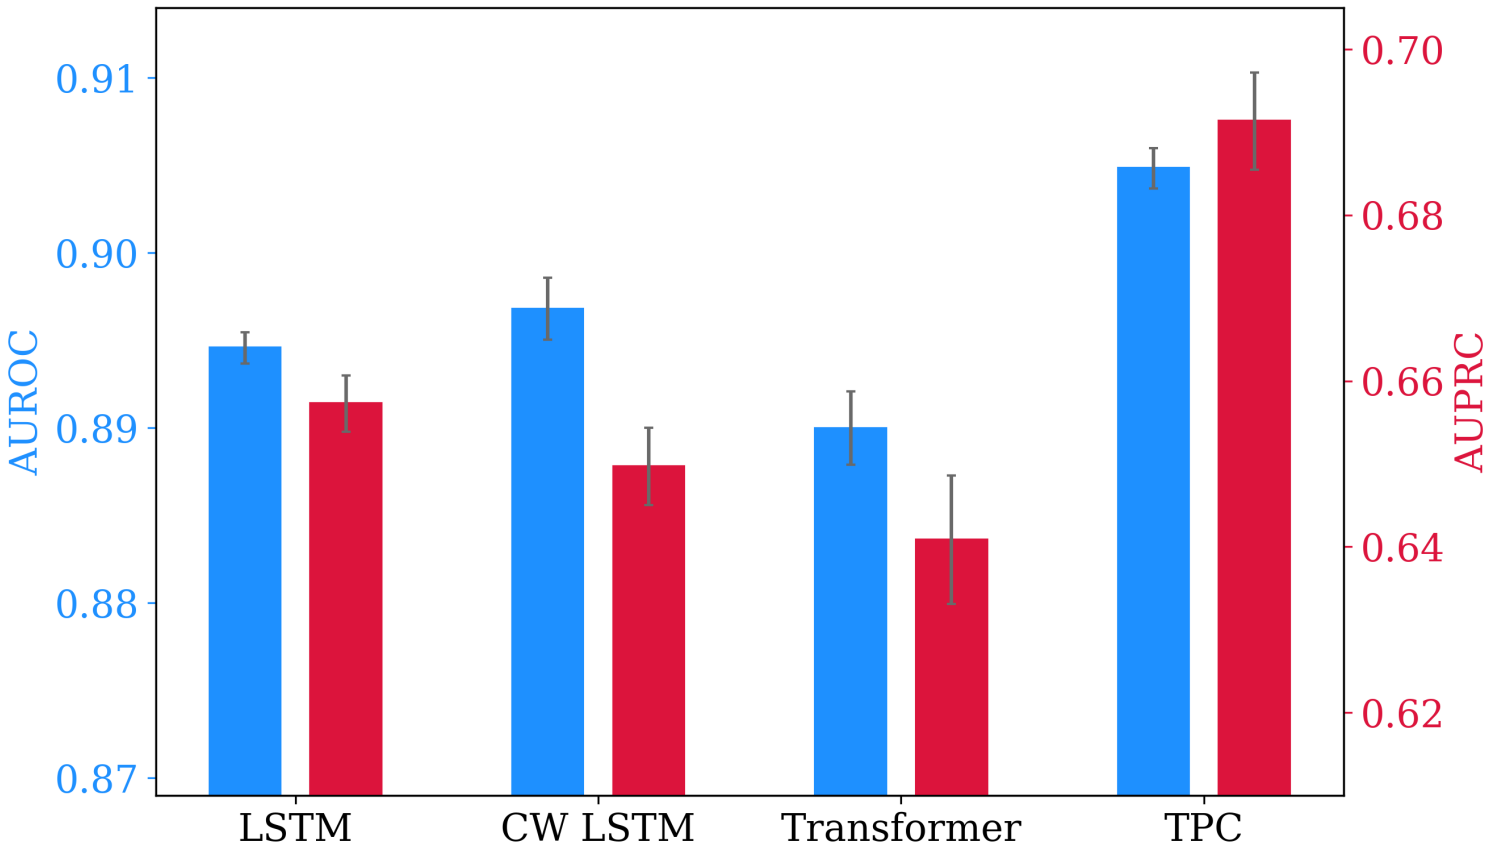

Ablation Study Results

- The temporal convolution performs better than the pointwise convolution
- However the best performance is gained when they are used together.
- The temporal convolution performance is helped by:
  - Skip connections.
  - Not sharing weights between features.

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.

Thank you!

To my funders:

- The Armstrong Fund
- The Frank Elmore Fund
- The Clinical School

...my supervisor and mentor:

Prof. Pietro Liò and Dr Stephanie Hyland