Sequential sub-network routing for multi-task learning on electronic health records

Subhrajit Roy, Diana Mincu, Eric Loreaux, Anne Mottram, Ivan Protsyuk, Natalie Harris, Sebastien Baur, Yuan Xue, Jessica Schrouff, Hugh Montgomery, Ali Connell, Nenad Tomasev, Alan Karthikesalingam, Martin Seneviratne

Google DeepMind

Introduction

Applications of machine learning (ML) to electronic health records (EHRs) have shown promise in predicting patient deterioration. One shortcoming is that the majority of these ML models are based on single-task learning (ST). By contrast, a clinician's mental model is multi-dimensional and takes into account a variety of competing risks [1].

Multi-task learning (MTL), which better reflects this multi-dimensional clinician mindset, has shown promise in improving model performance and training efficiency; however it often suffers from negative transfer - impaired learning if tasks are not appropriately selected [2].

We propose a sequential deep MTL architecture called sequential sub-network routing (SeqSNR) [3], which uses flexible parameter sharing and modularisation to control information sharing across multiple tasks thereby reducing the risk of negative transfer.

Dataset and prediction tasks

- **MIMIC-III Clinical Database**: An openly accessible database containing data from 36,498 adult patients across 52,038 admissions to critical care units at the Beth Israel Deaconess Medical Center between 2001 and 2012.
- **Prediction endpoints**: A suite of adverse events, interventions, and administrative endpoints: Acute Kidney Injury (AKI), continuous renal replacement therapy (CRRT) dialysis, vasoactive medications, first mechanical ventilation (MV), mortality, and remaining length of stay.

Methods

A high-level overview of the Sequential sub-network (SeqSNR) routing architecture.

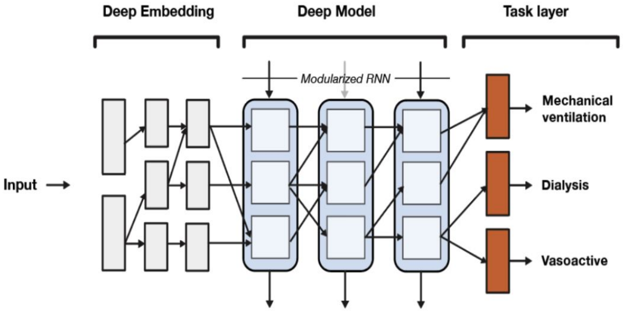

Results: Discriminatory performance

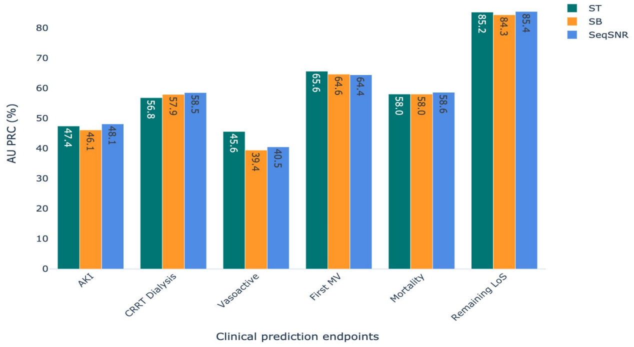

- SeqSNR > SB on 5 out of 6 tasks

- SeqSNR > ST on 4 out of 6 tasks

Results: Label efficiency

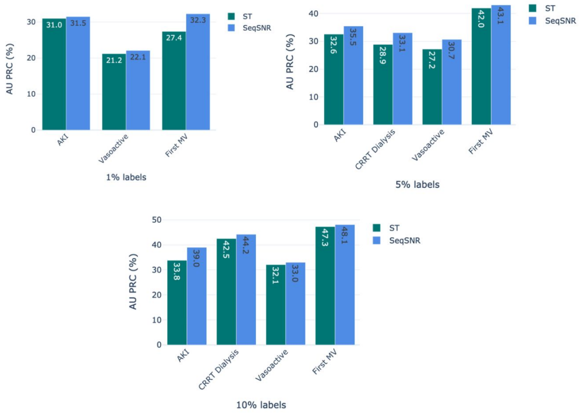

- We formulated prediction tasks with only a portion of the labels available for the primary prediction task (AKI, vasoactive, first MV)), but kept the entire dataset for the "helper tasks" (labs & vitals, mortality, remaining Los).

- SeqSNR outperformed ST across all tasks and all training data reduction percentages in a statistically significant way.

- Given the difficulties of annotating endpoint labels in EHR datasets, which frequently necessitates human evaluation by doctors, the ability to use numerous endpoints, could lessen the need for manual curation on more difficult endpoints.

References

Z. Zador (2019), "Multimorbidity states associated with higher mortality rates in organ dysfunction and sepsis: A data-driven analysis in critical care," Critical Care

M. McDermott (2021), "A Comprehensive Evaluation of Multi-task Learning and Multi-task Pre-training on EHR Time-series Data," ACM CHIL.

S. Roy (2021), "Multitask prediction of organ dysfunction in the intensive care unit using sequentials ubnetwork routing,"Journal of the American Medical Informatics Association.