Multi-task Cox Proportional Hazard Model for Predicting Risk of Unplanned Hospital Readmission

Megan Grzyb, Amber Zhang, Cristina Good, Khaled Khalil, Bochen Guo, Lu Tian, Jose Valdez, and Quanquan Gu
University of Virginia

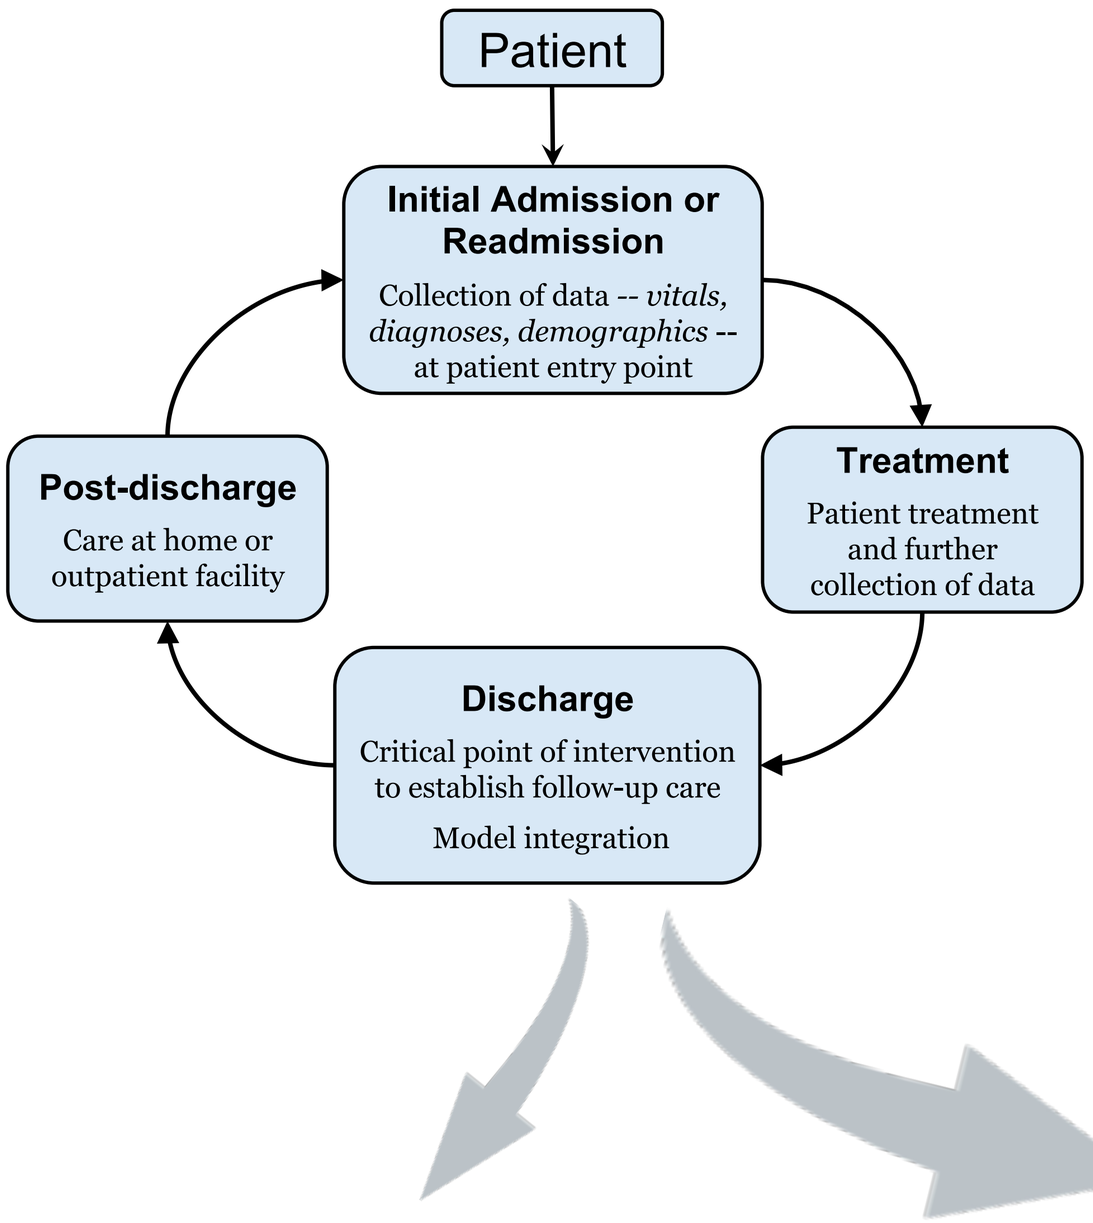

Unplanned Hospital Readmission

Readmission -- an urgent medical situation requiring hospitalization within 30 days of previous discharge

Costs $15 billion annually to care for readmissions

Medicare and Medicaid patients have highest likelihood of readmission at 18%

Penalizes hospitals if their readmission rates are higher than the national average, 15.2%, for Medicare and Medicaid patients

University of Virginia Medical Center

UVAMC estimates $764,000 of CMS penalties this year, due to 16.8% readmission rate

Data analytics department for UVAMC developed a Random Survival Forest (RSF) model to predict hospital readmissions

Transitional care facility partnering with UVAMC

Uses RSF output to provide targeted care to discharged patients

Cox Proportional Hazard Model

How it works: Semi-parametric model for survival analysis, learned by optimizing a partial likelihood function.

Assumptions:

- Hazard functions for different learning tasks, which calculate the rate of time before an event occurs, are proportional over time. Thus, patients in the same 30-day readmission window have the same baseline function, but with different attributes, to calculate time before readmission.
- No pre-determined distribution for failure time, which is when a patient experiences a readmission

Formulation:

$$l(B) = \sum_{i=1}^{D} B^T s_i - d_i \log\left[\sum_{j \in R_i} \exp\left(B^T z_j\right)\right]$$

Advantages:

- Allows sharper patient-by-patient predictions and better model interpretability than the current RSF model, through ability to see individualized patient readmission risk projections
- Focuses primarily on how the covariates affect the base hazard function.

Limitations:

- Assumes that the survival curves of different patient groups are naturally similarly shaped, which is not a practical assumption

Results: Concordance index of 0.70, after 10-fold cross validation.

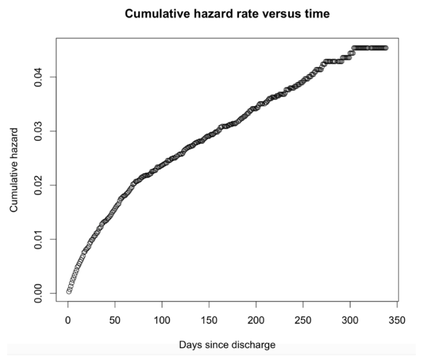

Cumulative Hazard Rate
The cumulative hazard rate is the accumulation of hazard over time. The figure on the left illustrates the increasing trend of the cumulative hazard rate for all patients in the same task.

Multi-task Cox Proportional Hazard Model

How it works: Semi-parametric model, learned by grouping patients into tasks by their diagnoses and finding a shared representation of significant features, to decrease the prediction error of each task.

Creating tasks: Patients grouped into task groups based on their admission diagnosis

Task	Diagnosis Code	Patient No.	Attributes
1	A	1	Age, income…
		2	Age, income…
		3	Age, income…
2	B	4	Age, income…
		6	Age, income…
…	…	…	…
N	Z	12	Age, income…

Model Fitting: Learning enhanced by finding attributes significant across many tasks

Assumptions:

- Covariates are generally significant or insignificant across all tasks
- Patient groupings by diagnoses are representative of inherent differences between the different tasks

Formulation:

$$l(B) = \sum_{k=1}^{K} \sum_{i=1}^{D} B_k^T s_{i,k} - d_{i,k} \log\left[\sum_{j \in R_{i,k}} \exp\left(B_k^T z_j\right)\right] + \lambda \|B\|_{2,1}$$

Advantages:

- Captures dependencies in outcomes at various timepoints by evaluating shared attributes across different tasks
- Creates a separate baseline hazard function for each task-grouping of diagnoses, to allow for more personalized predictions

Limitations:

- $L_{2,1}$norm seeks features that are either completely "relevant" or "irrelevant" across the different tasks.
- Questions related to how tasks are grouped, including number of observations within each group and the number of overall groups especially within the dataset used to train the multi-task model

Results: Concordance index of 0.52.

Conclusions and Future Works

Conclusion: *The Cox model outperforms the multi-task Cox model based on concordance index, because:*

- The sparse nature of the training dataset limited some of the task groupings to a small number of observations which in turn decreased the robustness of the model for certain tasks.
- Grouping tasks by diagnosis code might not be most suitable for multi-task modeling due to the nature of the patient population and how the diagnosis code feature is distributed.

Future Works: *The multi-task model is an innovative approach to predicting hospital readmissions, but can be improved by:*

- Larger, more robust dataset that is better suited for multi-task learning
- Greater choice and experimentation of tasks by which patients are grouped: utilizing age groups as an alternative to diagnosis codes for task groupings