Predicting Emergency Room (ER) Readmissions within 72 Hours of Discharge

Maeva Fincker, Sara Altman, and Hershel Mehta

Problem Statement

Problem

We used a Electronic Health Records (EHR) dataset to predict if a patient will return to the ER within 72 hours of discharge. A return within 72 hours of discharge suggests that a patient should not have been discharged previously. Hospitals could therefore use our models when considering whether to discharge a patient at a given time.

Data

Our data came from Stanford Hospital patients in the Stride-7 dataset. Each row in the dataset represents a hospital visit by a patient. We used the diagnosis, medication, and procedure codes associated with each visit as features, as well as demographic information collected about the patient. We used the patient's age at admission and discharge to label visits as followed by another visit within 72 hours of discharge or not and included 9 negative examples for each positive example. For each visit, we used only EHR data from the previous 12 months to predict the outcome.

Dense Model ("baseline")

Dense Model Data

STRIDE-7 dataset, which includes patient EHR data from Stanford Hospital Center (SHC) and Lucile Packard Children Hospital (LPCH).

Each patient's EHR data consists of a set of chronological visits to the hospital, where the following data is collected: medication codes, procedure codes , and diagnostic codes. See example table below.

patient_id	visit_id	Code 1	Code 2	Code 3	...	Code N
111	1	0	0	3		0
	2	0	0	0		1
222	1	2	0	0		0

Dense Model Overview (2 - 6 Layers)

Input	Network	Ouput
6 months EHR data prior to discharge	2 - 6 fully connected layers (layer sizes ranged from 5 - 512 neurons)	Probability of readmission within 72 hours

Dense Model Results

F1-score (Positive Only)

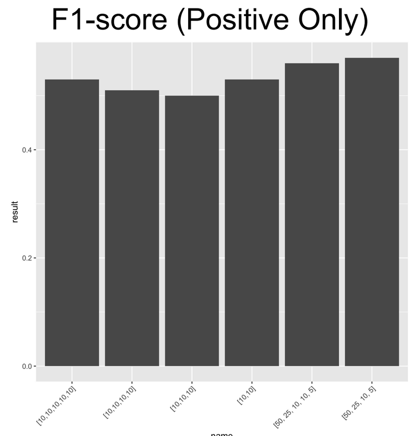

Error (1 - Accuracy)

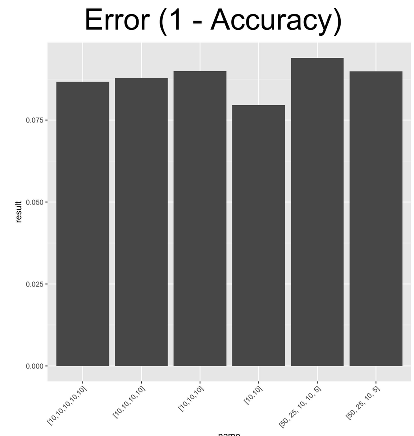

RNNs

RNN Data

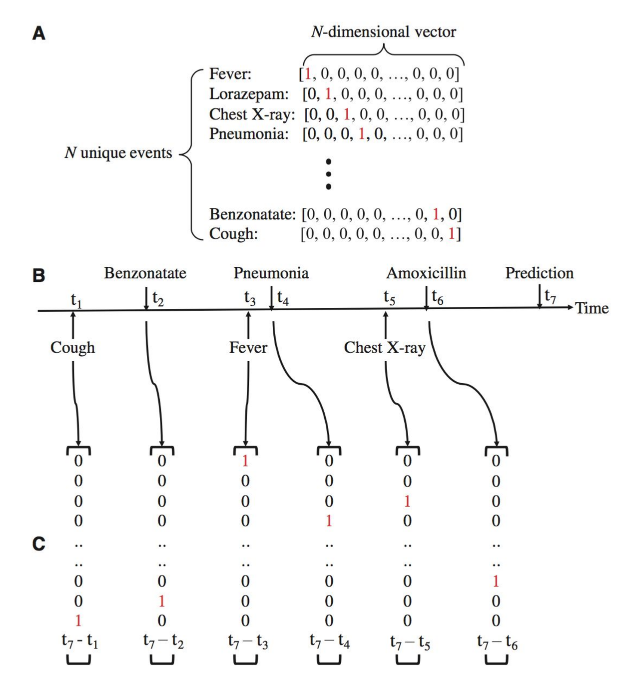

Adapted from 1. Choi E, Schuetz A, Stewart WF, Sun J. 2017. Using recurrent neural network models for early detection of heart failure onset. *J Am Med Inform Assoc.* 24(2):361–70

- Dataset filtered to only contain the 10,000 most common codes*.
- Each example is composed of a chronological code sequence comprising the codes from the 6 months before discharge.
- Each sequence is padded to the same max_length and fed to the network.

RNN Model

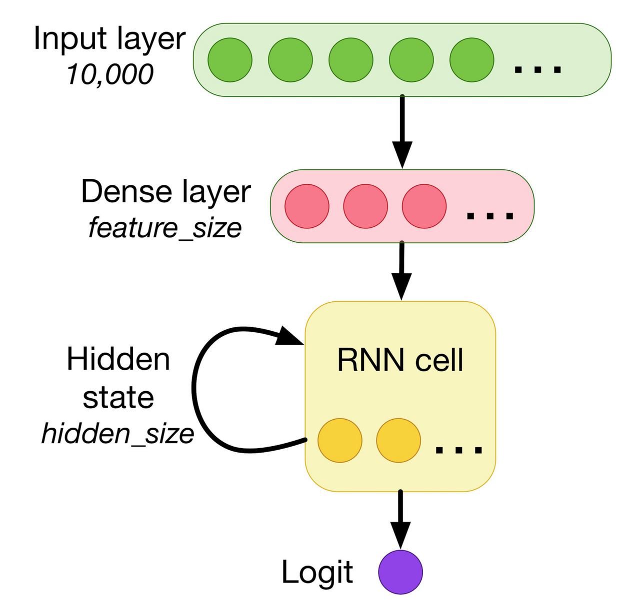

Loss: weighted logistic loss
Optimizier: Adam

Observations

- Training time is VERY long and depends highly on the length of the sequence.
- Weight initialization is critical.
- One-hot vectors are too sparse and carry too little information.
- Class imbalance can be mitigated by down-sampling negative examples / using a weighted loss.

Hyperparmeter search

learning rate	hidden_size	feature_size	max_length	model
0.1	10	100	10	gru
0.001	50	500	50	rnn
	100	1000	100	ensemble
			500	

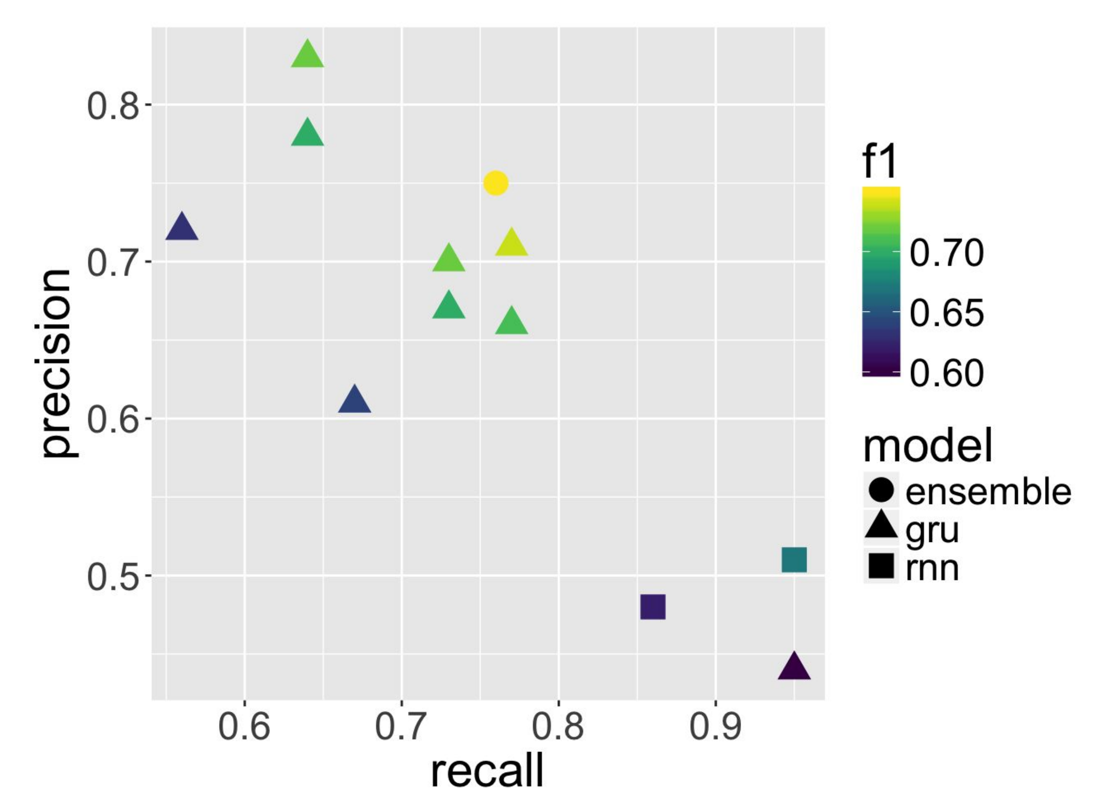

Model comparison

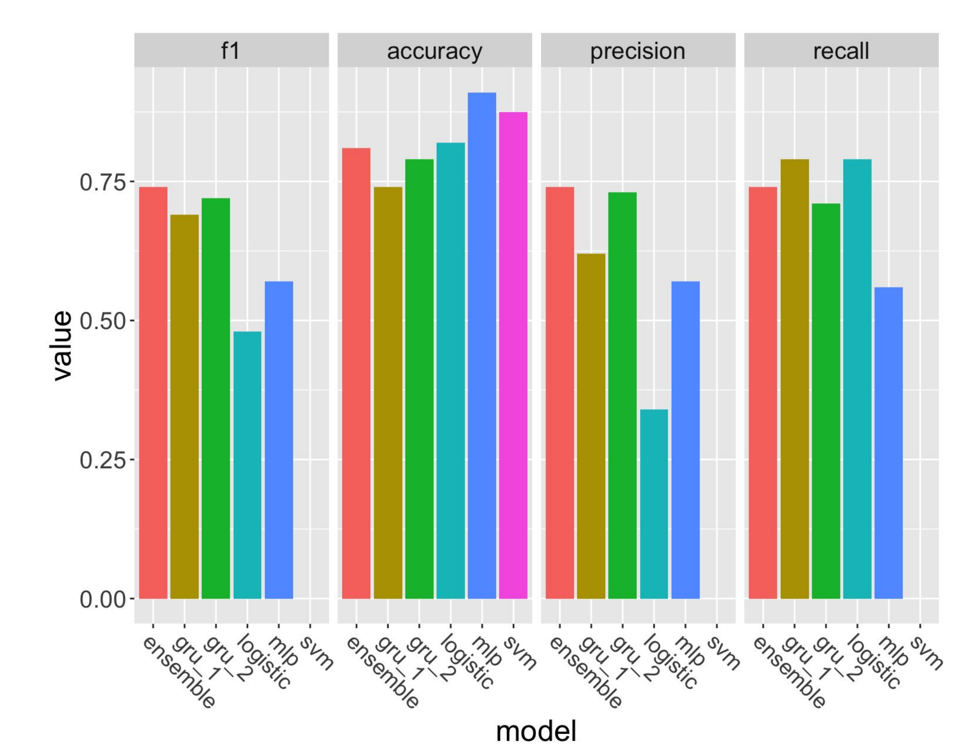

Conclusions

Discussion

Our problem proved very interesting, and our results are potentially clinically useful. Currently, there are no real metrics used to predict readmissions to the ER within 72-hours, which are extremely costly to hospitals, both in terms of finances and other resources (hospital beds, etc.). In the future, we would love to see our model implemented in a pilot program at the Stanford Hospital.

Code at: https://github.com/mfincker/ehr_deeplearning
Contact: mfincker@stanford.edu

Future Directions

- **Utilize Lab Results**: currently, we're only using diagnosis, procedure, and medication codes, but lab results would also be potentially helpful in this context
- **Medical Embeddings**: Many of our features represent higher level medical concepts, so it would be useful to collapse features into a lower dimensional "medical concept" space
- **Deeper Networks**: Currently we're working within memory space constrains, but our preliminary results suggest that even deeper networks could help reduce error resulting from Bias