Modeling Platelet Transfusion for The Stanford Blood Center: Inference Using Sentiment Analysis and Recurrent Neural Networks

Emily Guthrie, Rebecca Pattichis, and Rogelio Salinas DelaTorre
{eguthrie, pattichi, salin074}@stanford.edu

STANFORD BLOOD CENTER

Abstract

Platelets are a blood product that expire within 3 days of arriving to the hospital. The Stanford Hospital system wastes about 10% of platelets annually.

Researchers previously used aggregated data in order to predict usage, create a three-day ordering strategy, and thus reduce wastage. However, this ordering strategy was not implemented due to lack of human trust in models.

New research attempts to address this issue by using patient-level prediction. This project aims to aid this research by predicting which surgeries will need a platelet transfusion.

The two methods used for prediction are stochastic gradient descent on bag-of-words features and Recurrent Neural Networks.

Data

We used 146399 surgeries, encoded as follows:

Dept.	Rm.	Urgency	Proc1	Proc2	Proc3	*n* plt day of	*n* plt day after
cardiac	or	emergency	aortic dissection repair	valve replacement		3	2
otolaryngology	asc	elective	diverticulectomy endoscopic	zenkers	esophagoscopy	1	0
gastroenterology	endo	urgent	colonoscopy			0	0

Each surgery was classified as either:

+1 = needs platelet transfusion

−1 = no platelet transfusion

Method for Error Analysis

- Relatively few platelet transfusions (<3% of surgeries)
- Used precision and recall

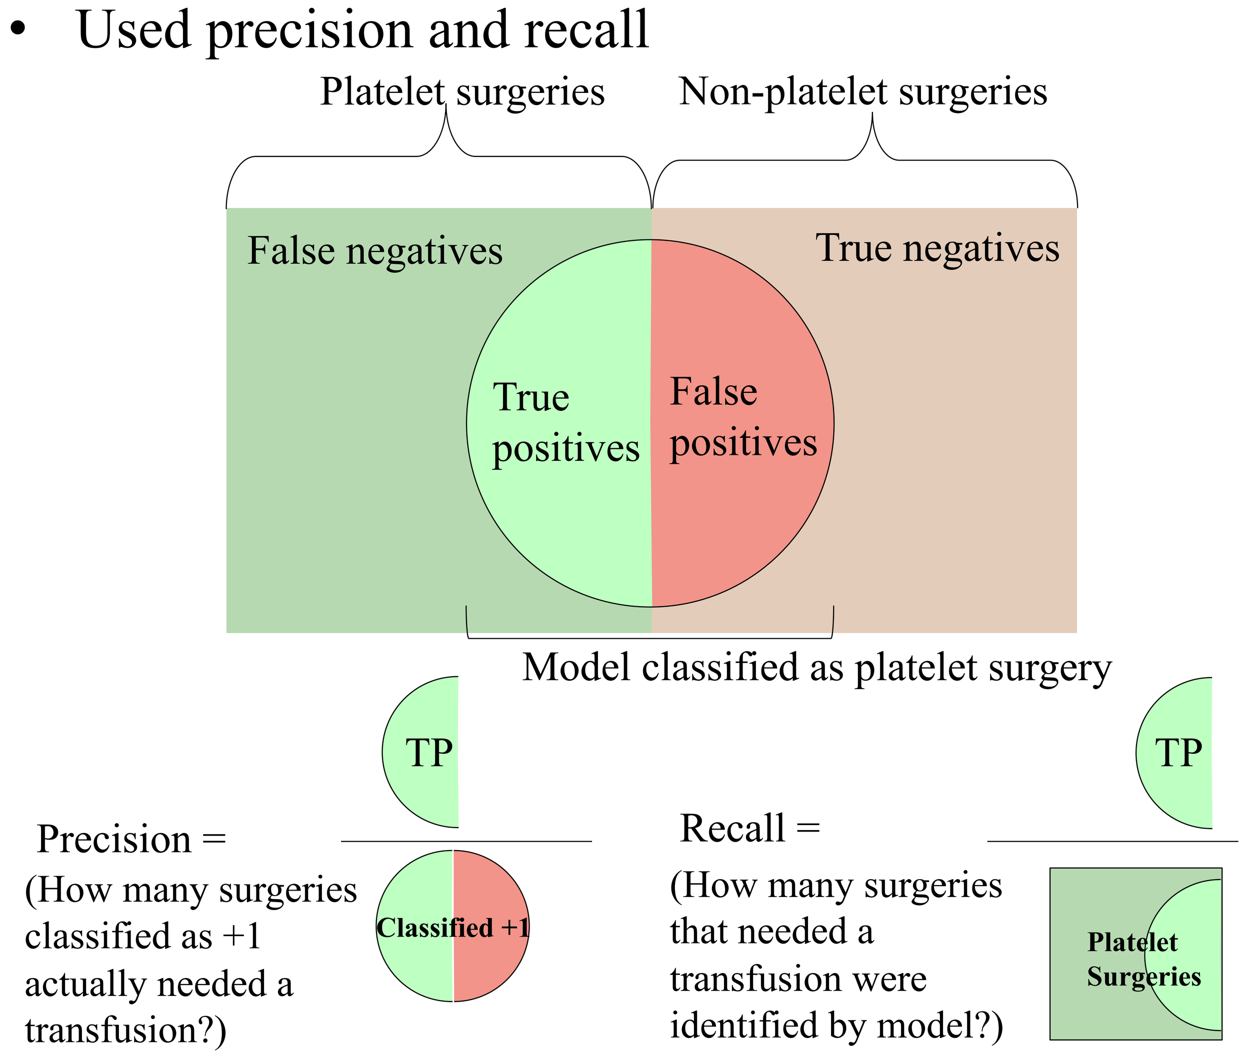

Modelling with Bag-of-Words and Stochastic Gradient Descent

- Begin with simple bag-of-words embedding
- Train via stochastic gradient descent, 100 epochs, $\eta = .001$
- Use 8-fold CV to estimate precision and recall
- Iterate on model as follows:

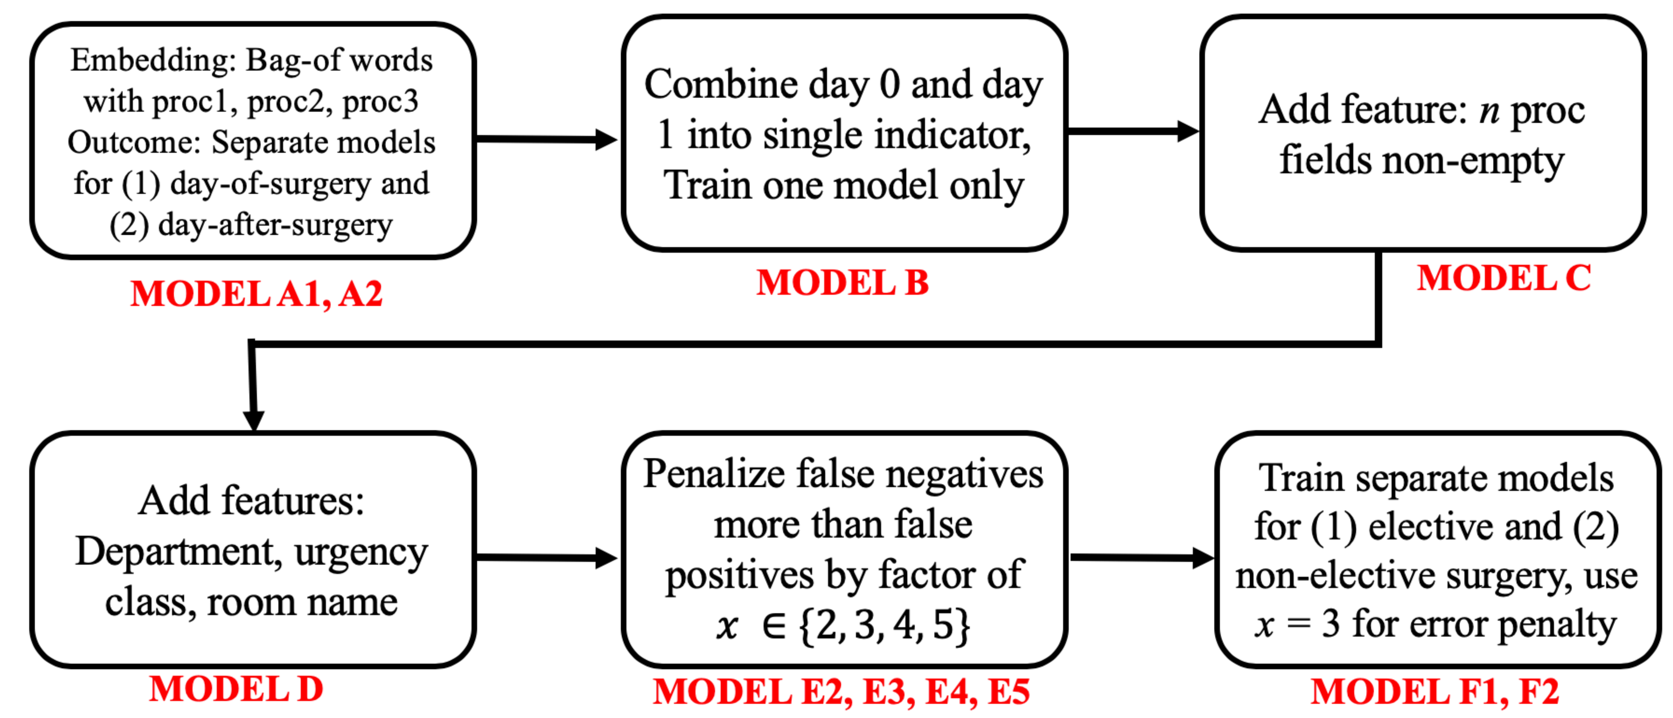

Figure 1: Bag-of-Words Results

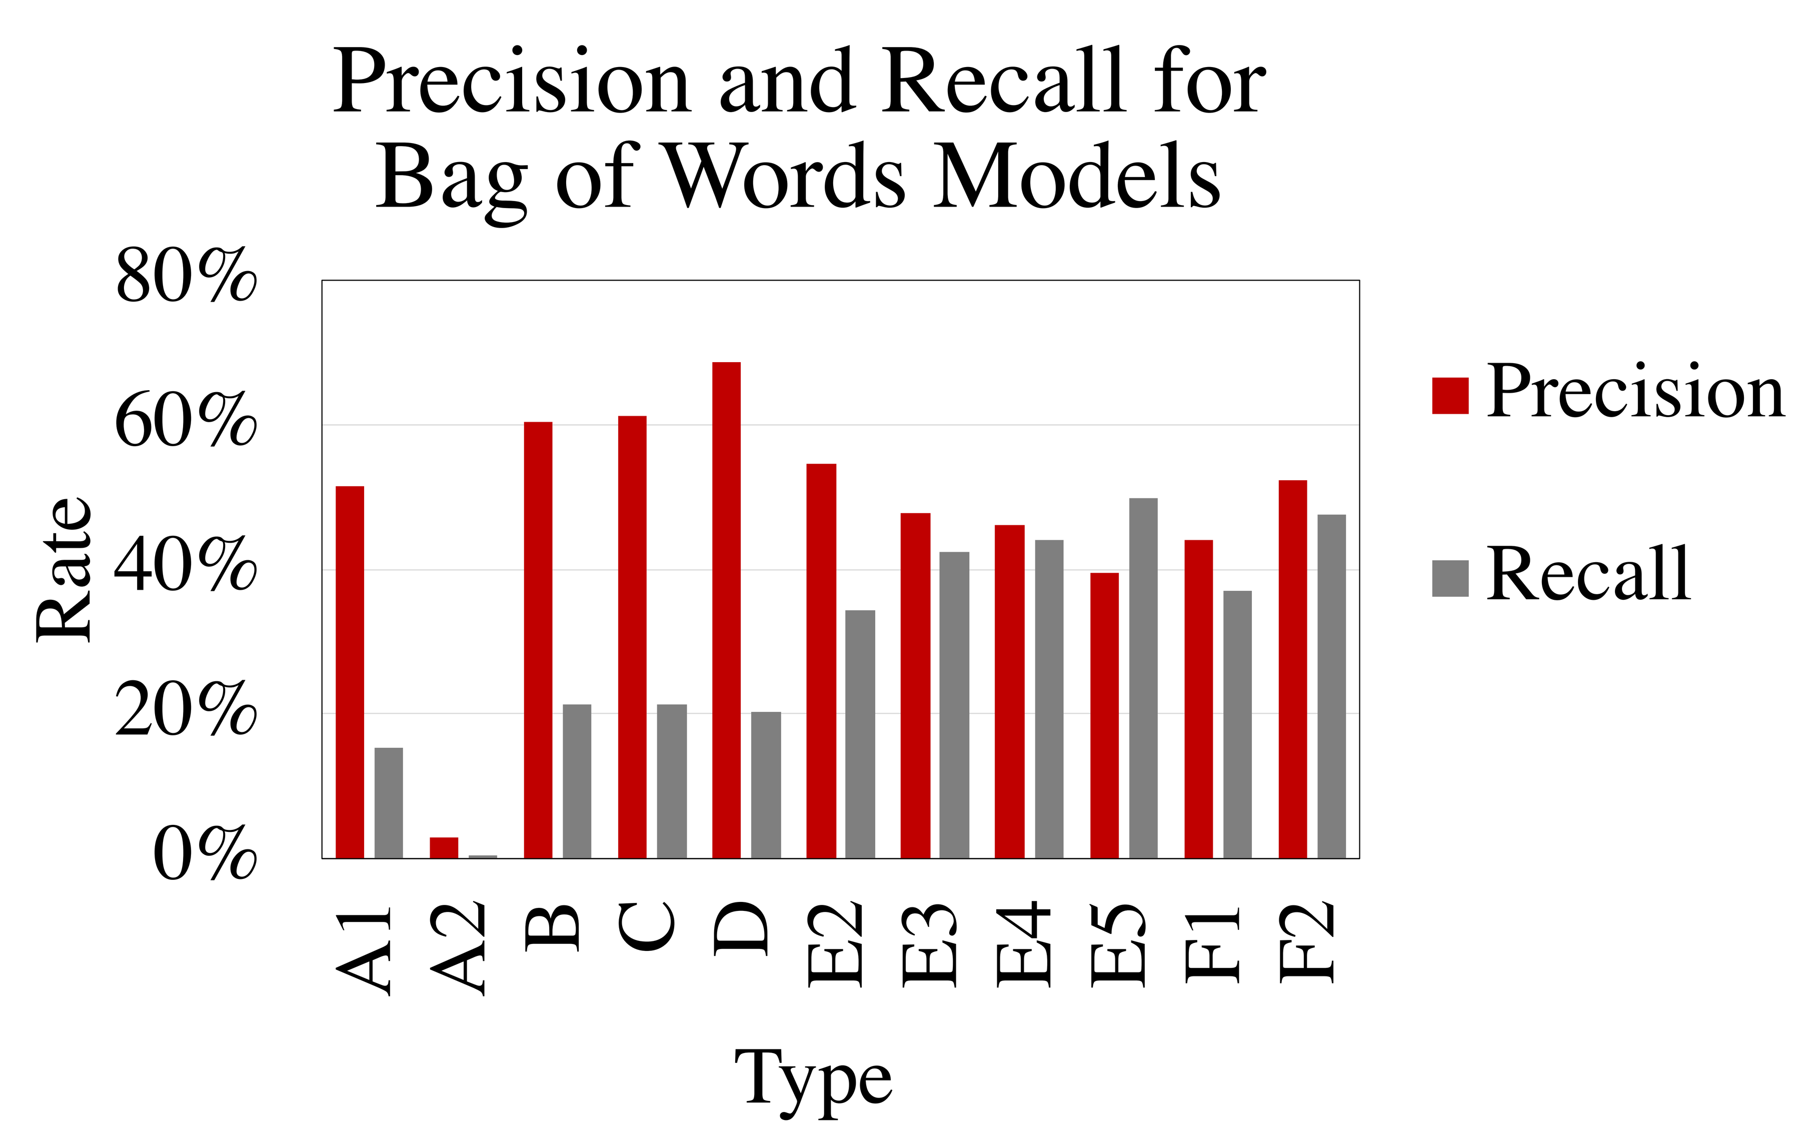

Model Take-aways:

- Prefer higher recall, since false negatives are worse than false positives
- Still need to test different error penalties for F1 and F2

Acknowledgements

Thank you to Dr. Tho Pham of the Stanford Blood Center for providing access to the data and support for the project. Thank you to Dr. Balasubramanian Narasimhan for advice on approaching this project. Thank you to Emily Flynn for helping with data access and logistics.

Modelling with Recurrent Neural Networks (GRU)

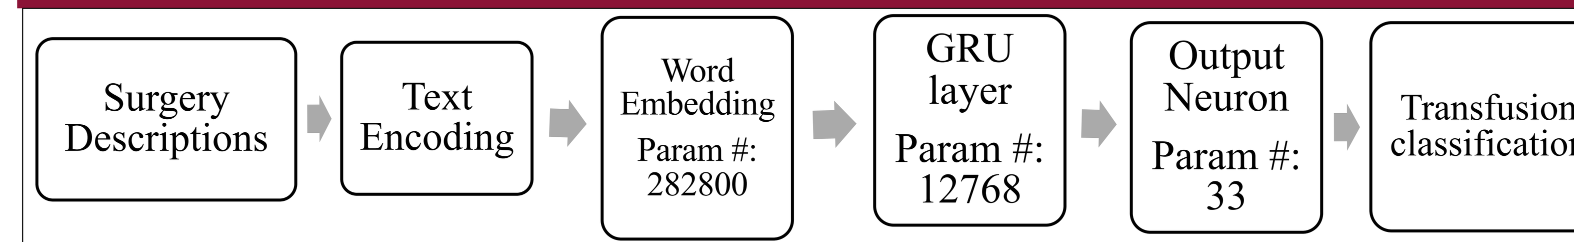

1. Model Layers:
   - Embedding: Maps each word to dense vector representations
   - GRU: 32 units, 0.2 dropout to avoid overfitting
   - Output Neuron: sigmoid activation function
2. Train: 25 epochs with train and validation data

Figure 2: RNN Results/Analysis

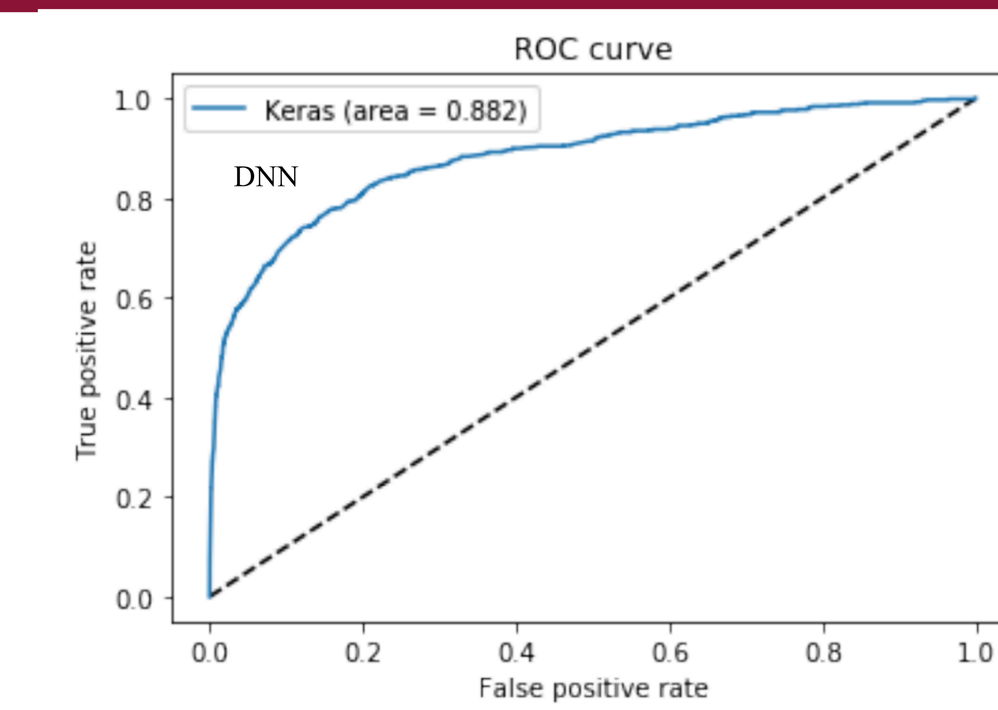

	Loss	Accuracy
Training set	.0686	97.87%
Validation set	.0915	97.56%

Achieved the following precision/recall values with a threshold = 0.5:

1. Precision: 65.96%
2. Recall: 27.27%

ROC curve highlights the model's classification abilities when threshold value is varied. The current AUC value is 0.882.

Future Improvements:

- Modify loss function: penalize false negatives more than false positives
- Implementing gloVe vectors to feed into RNN model
- Optimizing hyper parameters: test different models and choose the one with lowest validation error
- Deeper and/or more complex network (currently only employs one GRU layer)
- Choose threshold value based on the ROC curve computed on the training data
- Get final error metrics from true test set

References

A. Geron, *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow*, 2nd Edition, 2019.