




Delaney Sullivan*, Jonathan Wang*, Adam Wells*, Alex Wells*, Jonathan Chen MD/PhD
[dsull, jonwang1, adamw356, awells2, jonc101]@stanford.edu

Department of Medicine, Center for Biomedical Informatics Research, Stanford University School of Medicine, Stanford, CA *All authors contributed equally to this study.

Abstract

- The majority of clinical decisions lack evidence-based support, due to both the difficulty to perform randomized controlled patient experiments as well as the compliance of evidence-based guidelines ranging widely from 20-80% [1]; only about 11% of recommendation guidelines are backed by high quality evidence [2].
- Our group has developed a sequence prediction system to provide automated clinical decision support and predict patient progression through LSTM and feed-forward neural network models.
- We showcase two models, one to replace existing human derived order sets (f1: .492), and another to provide general patient timeline prediction (f1: .340)

Figure 1. ClinicNet Flow Diagram



Data

- Dataset spanning from 2008-2014 and including 2.2 million visits to Stanford Hospital was extracted from the Stanford Medicine Research Data Repository with IRB approval from HealthRx Laboratory.
- Each observation in our dataset corresponds to when a clinical item is ordered for a patient as recorded in the EHR, and includes specialty teams, labs, medications, images, nursing orders and other relevant columns that describe patient's current status.
- The "truth" for each observation is labeled as all the events that occur 24 hours after the clinical item was ordered

Features

- 7648 features encode information about what clinical items the patient received, medical teams he or she saw, and lab test results in the recent past. These features attempt to capture the patient status and context at the time of prediction task.
- Our labs, medications, nursing orders, etc. are all raw inputs. We include derived binary features for 1,7,30 and any day before for each of our clinical items.

Network Architecture

Feed-Forward Model

- The best performing feed-forward model had 5 hidden layers, each utilizing batch normalization and a ReLU activation function.
- Each output from the final hidden layer is passed through a sigmoid activation function, resulting in a 1452 x 1 dimensional output, corresponding to each of the 1452 clinical items.
- To prevent overfitting, Dropout (.05) was used during training.




Figure 2. Feed-Forward Architecture

$$Cost = \frac{-1}{m} \sum_{i=1}^m (q * y * \log a + (1 - y) * \log (1 - a))$$

Recurrent Neural Network Model

- The top performing Recurrent Neural Network model architecture used one layer of LSTM cells (800 hidden units/cell).
- Outputs were passed through a dense layer followed by a sigmoid activation to produce a 3003 x 1 dimensional output, corresponding to each of the 3003 patient progression outcomes.




Figure 3. Recurrent Neural Network Architecture

Equation 1. The weighted cross-entropy loss function for both models

Results

- Our RNN was trained to predict patient timeline events. We extracted 28,859 distinct 48 hour patient sequences and used a 90/5/5 training, dev, and test split. At each time-step, the RNN makes a patient progression prediction.
- We split 750,000 events into a 92/4/4 training, dev, and test split for our feedforward network, trained to generate order sets.

Table 1. LSTM and "Majority" Predictor Results

MODEL	Precision	Recall	F1	AUROC
LSTM	.366	.318	.340	.907
Majority	.140	.535	.221	.753

Figure 4. ClinicLSTM vs Majority


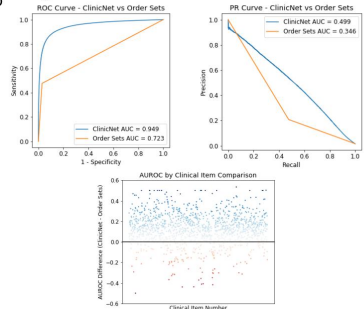


Table 2. Feed-Forward and Order Set Results

MODEL	Precision	Recall	F1	AUROC
Feed-Forward	.411	.612	.492	.949
Order Sets	.208	.476	.289	.723

Figure 5. ClinicNet Feed-Forward vs Order Sets



Discussion

- Our feed-forward model provides effective clinical decision support for physicians in a scalable and data-driven manner.
- We outperform existing, highly curated human-derived order sets, which are time consuming and unscalable, as well as all other previous models in the literature at predicting what clinical items a physician should order in the next 24 hours.
- Evidently, our feed-forward network architecture has the ability to pick up on patient context quite reliably.
- Using the data matrix we have constructed, we are also able to predict general timeline events, or patient progression, using a Recurrent Neural Network model that performs similarly to the f1, precision, and recall of existing human-derived order sets, but with more diversity, versatility and scalability in predictions.
- These models have the potential to be directly implemented into the clinical flow, supporting more effective, evidence-based practice.

Next Steps

- Begin the process for implementation of such a system into the clinical flow.
- Improve the performance of our algorithm by collecting more data (like clinical notes) and searching for hyper parameters over a longer period of time.
- Experiment with network architecture through adding attention for our Recurrent Neural Network model as well as time-awareness into our Feed-Forward model
- Look at prediction tasks for more specific subcases that will generate high accuracy numbers and thus be more reliable in decision support.

Acknowledgements & References

Thank you to Santosh Balasubramanian for helping us with data pre-processing, and Guillaume Genthial for network architecture guidance.

- Baker, A. (2001). Crossing the quality chasm: A new health system for the 21st century. *BMJ : British Medical Journal*, 323(7322), 1192. <http://dx.doi.org/10.1136/bmj.323.7322.1192>
- Tricci P, Allen JM, Kramer JM, Calif RM, Smith SC. Scientific Evidence Underlying the ACC/AHA Clinical Practice Guidelines. *JAMA*. 2009;301(8):831-841. doi:10.1001/jama.2009.205