Deep Learning Strategies for the Automatic Detection of Medication and Adverse Drug Events from Electronic Health Records

Susmitha Wunnava, MS, Xiao Qin, MS, Tabassum Kakar, MS,
Elke A. Rundensteiner, PhD and Xiangnan Kong, PhD
Worcester Polytechnic Institute, Worcester, MA, U.S.A

Background. Detecting the occurrence of adverse drug events (ADEs) and related medical information is an integral step towards the prevention of future critical ADE incidents threatening the public. Electronic health records (EHRs) of patients, a valuable source for potential ADE signals, are unstructured reports comprised of non-medical descriptive text and complex medical terminology[1]. Therefore, detecting a complete phrase that represents a named entity signaling an incidence is challenging. Recent research, leveraging popular deep learning approaches for the detection of medical information from EHRs, has outperformed state-of-the-art conditional random fields (CRF)[2]. An integration of recurrent neural networks (RNN) with CRF has also shown good performance[3].

Objective. Develop deep learning solution that achieves high accuracy in detection of medical entities from EHRs.

Methods. We have developed a named entity recognition (NER) system that utilizes the combined effectiveness of RNN (bi-directional long short-term memory (bi-LSTM)) and CRF by integrating them into one deep network architecture (Figure 1). Character-level representations are first obtained by a bi-LSTM over the sequence of characters in the input words. To represent words, a consolidated dense embedding, comprised of pre-trained medical word embeddings[3] concatenated with learned character-level embeddings, is used. The system is implemented using the Tensorflow framework.

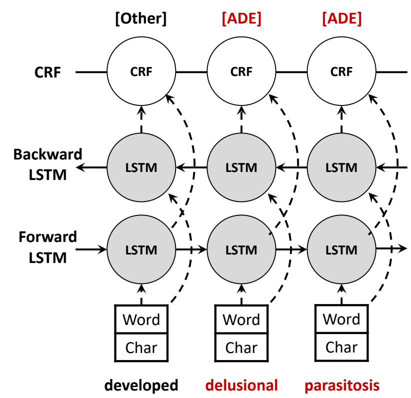

Figure 1: Deep learning network.

Findings. We evaluated our technology on de-identified EHR datasets released as part of the MADE 1.0 NLP challenge hosted by University of Massachusetts Medical School*. Our NER system is trained on 876 notes. It is evaluated on the test set of 213 notes. Our system ranked first for the Named Entity Recognition task in the MADE 1.0 challenge. Figure 2 shows our evaluation results for each of the entities and their overall micro-averaged F1-score. Our system with an overall F1-score of 0.84 outperformed the baseline systems 1, 2 and 3 with F1-score 0.79, 0.80, and 0.81, respectively.

Figure 2: Evaluation results on the final MADE 1.0 holdout test set.

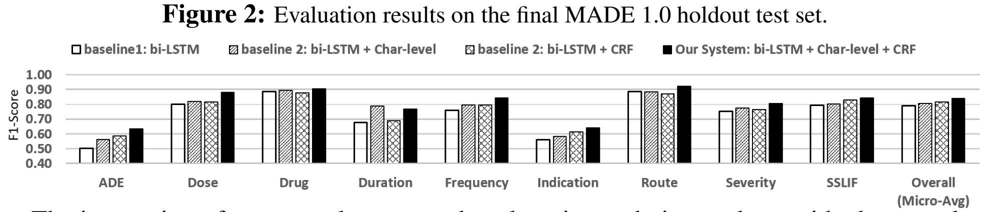

Conclusion. The integration of two complementary deep learning techniques along with character-level and word-level embedding in the input layer results in a deep learning architecture that achieves excellent extraction accuracy.

References

[1] Wunnava S, Qin X, Kakar T, Kong X, Rundensteiner EA, Sahoo SK, et al. One Size Does Not Fit All: An Ensemble Approach Towards Information Extraction from Adverse Drug Event Narratives. In: Proceedings of HEALTHINF; 2018. p. 176–188.

[2] Jagannatha AN, Yu H. Bidirectional RNN for Medical Event Detection in Electronic Health Records. In: Proceedings of NAACL HLT; 2016. p. 473–482.

[3] Jagannatha A, Yu H. Structured prediction models for RNN based sequence labeling in clinical text. In: Proceedings of EMNLP; 2016. p. 856–865.

*https://bio-nlp.org/index.php/projects/39-nlp-challenges