

A Neural Attention Model for Categorizing Patient Safety Events

Arman Cohan¹, Allan Fong², Nazli Goharian¹, and Raj Ratwani²

¹ Department of Computer Science, Georgetown University, Washington DC, USA
{arman,nazli}@ir.cs.georgetown.edu

² National Center for Human Factors in Healthcare, MedStar Health
Washington DC, USA
{allan.fong,raj.ratwani}@medicalhfe.org

Abstract. Medical errors are leading causes of death in the US and as such, prevention of these errors is paramount to promoting healthcare. Patient Safety Event reports are narratives describing potential adverse events to the patients and are important in identifying, and preventing medical errors. We present a neural network architecture for identifying the type of safety events which is the first step in understanding these narratives. Our proposed model is based on a soft neural attention model to improve the effectiveness of encoding long sequences. Empirical results on two large-scale real-world datasets of patient safety reports demonstrate the effectiveness of our method with significant improvements over existing methods.

Keywords: Natural Language Processing, Text Categorization, Medical Text Processing, Deep Learning

Introduction

There is an increasing demand for use of textual electronic health records and clinical notes to promote healthcare, and as such, In recent years NLP/IR have become increasingly important in understanding, searching, and analyzing medical information [22]. Human or system errors do occur frequently in the health centers, many of which can lead to serious harm to individuals. There are in fact an alarming number of annual death incidents (up to 200K) being reported due to medical errors [1]; medical errors are shown to be the third leading cause of death in the US [14]. Many healthcare centers have deployed patient safety event reporting systems to better identify, mitigate, and prevent errors [4]. Patient safety event reports are narratives describing a safety event and they belong to different safety categories such as “medication”, “diagnosis”, “treatment”, “lab”, etc. Recently, due to the importance of patient safety reports, more healthcare centers are enforcing patient safety reporting, resulting in an overwhelming number of daily produced reports. Manual processing of all these reports to identify important cases, trends, or system issues is extremely difficult, inefficient, and expensive. The first step in understanding and analyzing these events is to identify their general categories. This task is challenging because the

event descriptions can be very complex; the frontline staff usually focus more on taking care of the patient at the moment than to think through the classification schema when they later write a safety report. For example, an event where a patient fell after being given an incorrect medication might have been classified as “FALL” however, the fall could be due to a mis-medication and therefore belong to the “MEDICATION” safety event. Without the ability to correctly identify the medication category, such problems will not be addressed. Therefore, classifying the patient safety reports not only helps in further search and analytic tasks, but also it contributes to reducing the human reporting errors.

In this paper, we present a method for categorizing the Patient Safety Reports as the first step towards understanding adverse events and the way to prevent them. Traditional approaches of text categorization rely on sparse feature extraction from clinical narratives and then classifying the types of events based on these feature representations. In these conventional methods, complex lexical relations and long-term dependencies of the narratives are not captured. We propose a neural attention architecture for classifying safety events, which performs the feature extraction, and type classification jointly; our proposed architecture is based on a combination of Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) with soft attention mechanism. We evaluate our method on two large scale datasets obtained from two large healthcare providers. We demonstrate that our proposed method significantly improves over several traditional baselines, as well as more recent neural network based methods.

The proposed Neural Attention Architecture

Our proposed model for classifying patient safety reports is a neural architecture based on Convolutional Neural Networks (CNN) and Recurrent Neural Networks (RNN) utilizing a soft attention mechanism. Our architecture is partially similar to models by [12,11] in convolutional layers, to [19] in recurrent layer, and to [21] in the document modeling. Our point of departure is that unlike these works which are mainly targeted for sentence and short documents, we utilize a soft neural attention mechanism coupled with CNN and RNN to capture the more salient local features in longer sequences. Below we present the building blocks of our proposed architecture from bottom to the top.

Embedding layer. Given a sequence of words $S = \langle w_1; w_2; \dots; w_n \rangle$ where w_i are words in the sequence and “.” is the concatenation operation, the embedding layer represents S as an input vector $\mathbf{x} \in \mathbb{R}^{(m,d)}$ where d is the embedding dimension size and m is the maximum sequence length. x_i ’s can be either initialized randomly or by pre-trained word embeddings, and then they can be jointly trained with the model.

CNN. CNNs are feed-forward networks which include two main operations: *convolution* and *pooling*. Convolution is an operation on two functions (input and kernel) of real valued arguments [13]. In our context, in layer ℓ in the network, convolution operates on sliding windows of width k_ℓ on the input $\mathbf{x}_{\ell-1}$ and yields a feature map F_ℓ :

$$F_\ell^{(i)} = g(\mathbf{W}_\ell \cdot \mathbf{x}_{\ell-1}^{(i,k_\ell)} + \mathbf{b}_\ell) \quad (1)$$

where \mathbf{W}_ℓ and \mathbf{b}_ℓ are the shared weights and biases in layer ℓ , g is an activation function, and $\mathbf{x}^{(i,k_\ell)} = \langle x^{i-\frac{(k_\ell-1)}{2}}; \dots; x^{i+\frac{(k_\ell-1)}{2}} \rangle$ shows the sliding window of size k_ℓ centered at position i on the input. We use ReLU [5] for the activation function (In our experiments ReLU showed the best results among other activation functions). For pooling, we use “max-pooling” operation whose role is to down-sample the feature map and capture significant local features. Similar to [12], we use filters of sizes from 2 to 6 to capture local features of different granularities. The convolution layer allows the model to learn the salient features that are needed for identifying the type of the safety events.

RNN. Unlike CNNs which are local feature encoders, RNNs can encode large windows of local features and capture long temporal dependencies. Given an input sequence $\mathbf{h} = (x_1, \dots, x_T)$ where each $x_t \in \mathbb{R}^d$ is an input word vector of dimension d at time step t , an RNN computes the hidden states $\mathbf{h} = (h_1, \dots, h_T)$ and outputs $\mathbf{y} = (y_1, \dots, y_T)$ according to the following equations [8]:

$$h_t = g(W^{(hh)}h_{t-1} + W^{(xh)}x_t + b_h) \quad y_t = W^{(hy)}h_t + b_y \quad (2)$$

where W shows the weight matrices for the corresponding input, b denotes the biases, and g is the activation function. RNNs in theory, can capture temporal dependencies of any length. However, training RNNs in their basic form is problematic due to the *vanishing gradient* problem [16]. Long Short-Term Memory (LSTM) [10] is a type of RNN that has several gates controlling the flow of information to be preserved or forgotten, and mitigates the vanishing gradient problem. We use the LSTM formulation as in [9]. We also employ bidirectional LSTM to capture both forward and backward temporal dependencies. Using this layer, we capture the dependencies between local features along long sequences.

Neural attention. The trouble with RNNs for classification is that they encode the entire sequence into the vector at the last temporal step. While the application of RNNs have been successful in encoding sentences or short documents, in longer documents this can result in loss of information [3], and putting more focus on the recent temporal entries [18]. Bidirectional RNNs try to alleviate this problem by considering both the forward and backward context vectors. However, they suffer from the same problem in long sequences.

Inspired by work in machine-translation, to address this problem, we utilize the soft attention mechanism [7]. Neural attention allows the model to decide which parts of the sequence are more important instead of directly considering the context vector output by the RNN. Specifically, instead of considering the final cell state of LSTM for the classification, we allow the model to attend to the important timesteps and build a context vector c as follows:

$$c = \sum_{t=1}^T \alpha_t h_t \quad (3)$$

where α_t are weights computed at each timestep t for the state h_t and are computed as follows:

$$\alpha_t = \frac{\exp(e_t^\top z)}{\sum_{k=1}^T \exp(e_k^\top z)} \quad (4) \quad e_t = f_{\text{ATT}}(h_t) \quad (5)$$

Table 1: Dataset characteristics

	Number of Reports	Number of categories	Avg. length (char)	Stdev. length (char)
Dataset 1	82,281	20	410	321
Dataset 2	1,625,512	9	327	174

Table 2: Categories in the larger dataset (Dataset 2)

Category	Count	Category	Count
Procedure/Treatment/Test Error	370K	Miscellaneous	140K
Medication Error	135K	Adverse Drug Reaction	34K
Fall	242K	Equipment/Supplies/Devices	34K
Procedure/Treatment/Test Complication	233K	Transfusion	23K
Skin Integrity	234K		

where f_{ATT} is a function whose role is to capture the importance of h_{t_i} and z is a context vector that is learned jointly during training. We use a feed-forward network with “tanh” activation function for f_{ATT} . The context vector c is then fed to a fully-connected and then a softmax layer to perform final classification.

Experiments

Setup. We evaluate the effectiveness of our model on two large scale patient safety data obtained from a large healthcare providers in mid-Atlantic US and the Institute for Safe Medication Practices (ISMP). ISMP serves as a safe harbor for all PSE reports from hospitals in Pennsylvania, US. The dataset that was analyzed contains all categories of safety reports (fall, medication, surgery, etc.) and is not limited to medication reports. This study was approved by the Med-Star Health Research Institute Institutional Review Board (protocol 2014-101). The characteristics of the data and the categories are shown in tables 1 and 2. We split the data with stratified random sampling into 3 sets: train, validation, and test. We tune the parameters of the neural models on the validation set and the test set remains unseen to the models. We compare our results with conventional text classification models (bag of words feature representation with different types of classifiers), as well as related work on neural architectures (CNNs, RNNs and Bidirectional RNNs and their combinations). For space limitation, we do not explain the details of the baselines and refer the reader to the corresponding citations in Table 3. We report accuracy and average F1-score results for the categories which are standard evaluation metrics for this task.

Implementation. We used Keras and TensorFlow for the implementation. We empirically made the following design choices: We used Word2Vec [15] for training the embeddings on both general (Wikipedia) and domain specific corpora (PubMed), similar to [17]. We used dropout rates of 0.25 for the recurrent and 0.5 for the convolutional layers. Training was done using Adam optimizer with categorical cross entropy loss; we also applied early stopping for the training. Number of epochs for the larger dataset was 2 and for small dataset 6.

Table 3: Results on the each dataset on both the validation and test sets. Numbers are percentages. Last row shows our method. † (‡) shows statistically significant improvement (McNemar’s test) over the next best performance with $p < 0.05$ ($p < 0.01$).

Methods	Dataset 1				Dataset 2			
	Val		Test		Val		Test	
	Acc	F1	Acc	F1	Acc	F1	Acc	F1
SVM [20]	70.7	70.3	70.9	70.6	84.8	84.0	84.7	83.9
MNB [20]	71.2	71.5	71.0	72.3	79.2	79.9	79.0	79.6
XGB [2]	71.4	69.9	72.1	70.8	76.8	75.7	76.7	75.5
cBoW [23]	67.5	62.6	68.0	63.4	84.8	84.2	84.6	84.1
Adaptive cBoW [23]	69.2	63.4	70.6	69.6	83.9	84.3	84.8	84.8
CNN [12]	73.2	70.7	72.2	69.5	83.6	83.1	82.7	83.5
RNN [6]	76.0	74.6	74.5	72.9	84.0	84.2	83.8	83.2
Bi-RNN [6]	76.3	74.5	75.2	73.6	84.7	84.3	84.6	84.5
CNN-BiRNN [19]	77.8	76.9	76.6	76.4	89.3	85.9	86.8	84.6
Att-CNN-BiRNN (ours)	78.3 †	77.2	78.1 ‡	77.3 ‡	89.1	88.1 ‡	88.9 ‡	88.0 ‡

Results. Table 3 demonstrates our main results. As illustrated, our method (last row) significantly outperforms all other methods in virtually all the datasets. This shows the general effectiveness of our model in comparison with the prior work. We observe that our method’s performance improvement is slightly larger in the second (larger) dataset. This is expected since our model can better learn the parameters when trained on larger data. Improvement over RNN and CNN-Bi-RNN baselines shows the effectiveness of the neural soft attention in capturing salient parts of the sequence in comparison with the models without attention.

Error Analysis. While our method effectively outperforms the prior work, we conducted error analysis to better understand the cases that our method fails to correctly perform categorization. In particular, we observed that for both datasets, many wrongly classified samples in the categories were misclassified as the “MISCELLANEOUS” category. This pattern was more common for the categories with smaller training samples. This shows that the model learns a broader set of texts for the “MISCELLANEOUS” category, which is expected, given the broad nature of this category. We also observed some misclassified samples in the categories that are closely related together. For example in dataset 1, 32% of the misclassified samples in the “BLOOD-BANK” category were classified as “LAB/SPECIMEN”. A similar pattern was observed for the “DIAGNOSIS” and “MEDICATION” safety events. These closely related categories usually have




Fig. 1: Performance for each category based on its relative size to the dataset.

overlaps in terms of training data and this makes it hard for the model to differentiate the edge cases. We furthermore observe that the performance on each category correlates with the number of samples in that category. Figure 1 shows this correlation. We observe that generally, our method performs better with the categories of larger relative size. While the correlation is stronger for dataset 1, both datasets show similar trends. This shows that having more training samples helps our model in better learning the characteristics of that particular category and results in higher performance.

Conclusion

We presented a neural network model based on a soft attention mechanism for categorizing patient safety event reports. We demonstrated the effectiveness of our model on two large-scale real-world datasets and we obtained significant improvements over existing methods. The impact of our method and results is substantial on the patient safety and healthcare, as better categorization of events results in better patient management and prevention of harm to the individuals.

Acknowledgments

We thank the 3 anonymous reviewers for their helpful comments. This project was funded under contract/grant number Grant R01 HS023701-02 from the Agency for Healthcare Research and Quality (AHRQ), U.S. Department of Health and Human Services. The opinions expressed in this document are those of the authors and do not necessarily reflect the official position of AHRQ or the U.S. Department of Health and Human Services.

References

1. American Hospital Association: Fast facts on US hospitals (2013)
2. Chen, T., Guestrin, C.: Xgboost: A scalable tree boosting system. In: Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 785–794. ACM (2016)
3. Cho, K., van Merriënboer, B., Bahdanau, D., Bengio, Y.: On the Properties of Neural Machine Translation: EncoderDecoder Approaches. Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation pp. 103–111 (2014), <http://arxiv.org/abs/1409.1259>
4. Clarke, J.R.: How a system for reporting medical errors can and cannot improve patient safety. *The American surgeon* 72(11), 1088–91; discussion 1126–48 (2006), <http://www.ncbi.nlm.nih.gov/pubmed/17120952>
5. Dahl, G.E., Sainath, T.N., Hinton, G.E.: Improving deep neural networks for lvcsr using rectified linear units and dropout. In: Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on. pp. 8609–8613. IEEE (2013)
6. Dai, A.M., Le, Q.V.: Semi-supervised Sequence Learning. In: Neural Information Processing Systems (NIPS). pp. 3079–3087 (2015)

7. Dzmitry Bahdanau, Bahdanau, D., Cho, K., Bengio, Y.: Neural Machine Translation By Jointly Learning To Align and Translate. *Iclr 2015* pp. 1–15 (2014), <http://arxiv.org/abs/1409.0473v3>
8. Elman, J.L.: Finding structure in time. *Cognitive Science* 14(1 990), 179–211 (1990)
9. Graves, A., Jaitly, N.: Towards End-To-End Speech Recognition with Recurrent Neural Networks. In: *JMLR Workshop and Conference Proceedings*. vol. 32, pp. 1764–1772 (2014), <http://jmlr.org/proceedings/papers/v32/graves14.pdf>
10. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural computation* 9(8), 1735–1780 (1997)
11. Kalchbrenner, N., Grefenstette, E., Blunsom, P.: A Convolutional Neural Network for Modelling Sentences. In: *Association for Computational Linguistics*. pp. 655–665 (jun 2014)
12. Kim, Y.: Convolutional Neural Networks for Sentence Classification. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP 2014)*. pp. 1746–1751 (2014), <http://emnlp2014.org/papers/pdf/EMNLP2014181.pdf>
13. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86(11), 2278–2323 (1998)
14. Makary, M.A., Daniel, M.: Medical error the third leading cause of death in the US. *Bmj* 2139(May), 1–5 (2016)
15. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Distributed Representations of Words and Phrases and their Compositionality. In: *Neural Information Processing Systems (NIPS)*. pp. 1–9 (2013)
16. Pacanu, R., Mikolov, T., Bengio, Y.: On the Difficulties of Training Recurrent Neural Networks. *International Conference on Machine Learning (ICML) 28(2)*, 1310–1318 (2013)
17. Soldaini, L., Goharian, N.: Learning to Rank for Consumer Health Search: a Semantic Approach. In: *ECIR*. Springer (2017)
18. Sutskever, I., Vinyals, O., Le, Q.V.: Sequence to sequence learning with neural networks. In: *Neural Information Processing Systems (NIPS)*. pp. 1–9 (2014), <http://papers.nips.cc/paper/5346-sequence-to-sequence-learning-with-neural>
19. Tang, D., Qin, B., Liu, T.: Document Modeling with Gated Recurrent Neural Network for Sentiment Classification. In: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. pp. 1422–1432. No. September (2015), <http://aclweb.org/anthology/D15-1167>
20. Wang, S., Manning, C.: Baselines and Bigrams: Simple, Good Sentiment and Topic Classification. In: *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics*. pp. 90–94. No. July (2012)
21. Yang, Z., Yang, D., Dyer, C., He, X., Smola, A., Hovy, E.: Hierarchical Attention Networks for Document Classification. In: *Proceedings of NAACL-HLT 2016*. pp. 1480–1489 (2016), <https://www.aclweb.org/anthology/N/N16/N16-1174.pdf>
22. Yates, A., Goharian, N., Frieder, O.: Extracting Adverse Drug Reactions from Social Media. In: *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15)*. pp. 2460–2467. No. ii (2015)
23. Zhao, H., Lu, Z., Poupart, P.: Self-Adaptive Hierarchical Sentence Model. In: *IJ-CAI*. pp. 4069–4076. AAAI Press (2015), <http://arxiv.org/abs/1504.05070>