

# Hybrid Text Feature Modeling for Disease Group Prediction using Unstructured Physician Notes<sup>\*</sup>

Gokul S Krishnan<sup>[0000-0002-1344-4722]</sup> and Sowmya Kamath<sup>S[0000-0002-0888-7238]</sup>

Healthcare Analytics and Language Engineering (HALE) Lab,  
Department of Information Technology,  
National Institute of Technology Karnataka  
Surathkal, INDIA - 575025  
gsk1692@gmail.com, sowmyakamath@nitk.edu.in

**Abstract.** Existing Clinical Decision Support Systems (CDSSs) largely depend on the availability of structured patient data and Electronic Health Records (EHRs) to aid caregivers. However, in case of hospitals in developing countries, structured patient data formats are not widely adopted, where medical professionals still rely on clinical notes in the form of unstructured text. Such unstructured clinical notes recorded by medical personnel can also be a potential source of rich patient-specific information which can be leveraged to build CDSSs, even for hospitals in developing countries. If such unstructured clinical text can be used, the manual and time-consuming process of EHR generation will no longer be required, with huge person-hours and cost savings. In this article, we propose a generic ICD9 disease group prediction CDSS built on unstructured physician notes modeled using hybrid word embeddings. These word embeddings are used to train a deep neural network for effectively predicting ICD9 disease groups. Experimental evaluation showed that the proposed approach outperformed the state-of-the-art disease group prediction model built on structured EHRs by 15% in terms of AUROC and 40% in terms of AUPRC, thus proving our hypothesis and eliminating dependency on availability of structured patient data.

**Keywords:** Topic Modeling · Word Embedding · Natural Language Processing · Machine Learning · Healthcare Informatics · Mortality Prediction

## 1 Introduction

Electronic Health Records (EHRs) and predictive analytics have paved the way for Clinical Decision Support Systems (CDSSs) development towards realizing intelligent healthcare systems. The huge amount of patient data generated in

---

<sup>\*</sup> This work is funded by the Government of India's DST-SERB Early Career Research Grant (ECR/2017/001056) to the second author.

hospitals in the form of discharge summaries, clinical notes, diagnostic scans etc, present a wide array of opportunities to researchers and data scientists for developing effective methods to improve patient care. CDSSs built on techniques like data mining and machine learning (ML) have been prevalent in healthcare systems, aiding medical personnel to make informed, rational decisions and execute timely interventions for managing critical patients. Some existing ML based CDSSs include – mortality prediction [16, 23, 10, 21, 11, 5], hospital readmission prediction [12, 22, 14], length of stay prediction [11, 24, 1], generic disease or ICD9<sup>1</sup> group prediction [21, 19, 6] and so on. However, it is significant to note that, most existing CDSSs depend on structured patient data in the form of Electronic Health Records (EHRs), which favors developed countries due to large scale adoption of EHRs. However, healthcare personnel in developing countries continue to rely on clinical notes in the form of free and unstructured text, as structured EHR adoption rate is extremely low in these countries. Moreover, the conversion of patient information recorded in the form of text/clinical reports to a standard structured EHR is a labor-intensive, manual process that can result in loss of critical patient-specific information that might be present in the clinical notes. For this reason, there is a need for alternative approaches of developing CDSSs that do not rely on structured EHRs.

ICD9 disease coding is an important task in a hospital, during which a trained medical coder with domain knowledge assigns disease-specific, standardized codes called ICD9 codes to a patient’s admission record. As hospital billing and insurance claims are based on the assigned ICD9 codes, the coding task requires high precision, but is often prone to human error, which has resulted in an annual spending of \$25 billion in the US to improve coding efficacy [25, 7]. Hence, automated disease coding approaches have been developed as a solution to this problem. Currently, this is an area of active research [26, 17, 25, 3, 4, 2], but performance that has been recorded so far is below par indicating huge scope for improvement, thus making automated ICD9 coding an open research problem. Existing ICD9 coding methods make use of only discharge summaries for coding, similar to the process adopted by human medical coders. However, as the records are digitized, other textual records under the same admission identifier might be available, which may provide additional patient-specific insights pertaining to the diagnosed diseases. Further, for ICD9 disease coding to be effective, the correct determination of generic disease categories or groups prior to specific coding is very crucial as information regarding generic disease groups can be informative for specific ICD9 coding. With this in mind, we explore other textual sources of patient data, utilizing physician notes to predict the generic ICD9 disease groups for the patient, following which actual ICD9 disease code prediction can be achieved. In this article, we focus on effective prediction of ICD9 disease groups and the code prediction will be considered in future work.

While most existing works aim to predict ICD9 codes rather than ICD9 disease groups, Purushotham et al. [21] performed benchmarking of three CDSS

---

<sup>1</sup> ICD-9-CM: International Classification of Diseases, Ninth Revision, Clinical Modification

tasks - ICU mortality prediction, ICD9 group prediction and hospital readmission prediction on ICU patients data using Super Learner and Deep Learning techniques. The ICD9 group prediction task discussed in their work is a generic disease group prediction task that is a step prior to ICD9 coding, which can also be a viable disease risk prediction CDSS for caregivers. Their work uses feature subsets extracted from structured patient data, where the best performing model uses multitude of numerical features such as input events (volume of input fluids through IV), output events (quantity of output events like urinary and rectal), prescriptions and other ICU events like chart events (readings recorded in the ICU) and lab events (lab test values). As deep learning models can effectively learn from numerous raw feature values quite well, their approach achieved good results. However, in a real world scenario, to measure all these values and then using them as inputs to a prediction model for an outcome, will cause an inevitable and significant time delay, during which the patient condition can worsen. CDSS prediction models that use unstructured clinical text data can predict desired outcomes with lower latency and also help improve the prediction performance as these raw text sources contain more patient-specific information.

Some recent works [15, 9, 8] focussed on investigating the effectiveness of modeling various unstructured text reports such as radiology and nursing notes of ICU patients for disease group prediction. With the objective of continuing such analyses on other kinds of unstructured clinical reports, we intend to focus on physician notes in this work. Additionally, we also analyze how they can be used in predicting ICD9 disease groups by effectively modeling these physician notes. In this article, a hybrid feature modeling approach that uses hybrid clinical word embeddings to generate quality features which are used to train and build a deep neural network model to predict ICD9 disease groups is presented. We show the results of a benchmarking study of our proposed model (modeled on unstructured physician notes) against the current state-of-the-art model for disease group prediction (built on structured clinical data). The rest of this article is structured as follows: Section 2 describes the proposed approach in detail, followed by experimental results and discussion in Section 3. We conclude the article with prospective future work and research directions.

## 2 Materials & Methods

The proposed approach for disease group prediction is depicted in Fig. 1. We used physicians' clinical notes in unstructured text form from the MIMIC-III [13] dataset for our experiments. MIMIC-III contains data relating to 63,000 ICU admissions of more than 46,000 patients admitted in Beth Israel Hospital, New York, USA, during 2001-2012. We extracted only the physician notes from the 'NOTEVENTS' table, which resulted in a total of 141,624 physician notes generated during 8,983 admissions. As per MIMIC documentation, physicians have reported some identified errors in notes present in the 'NOTEVENTS' table. As these notes can affect the training negatively, records with physician

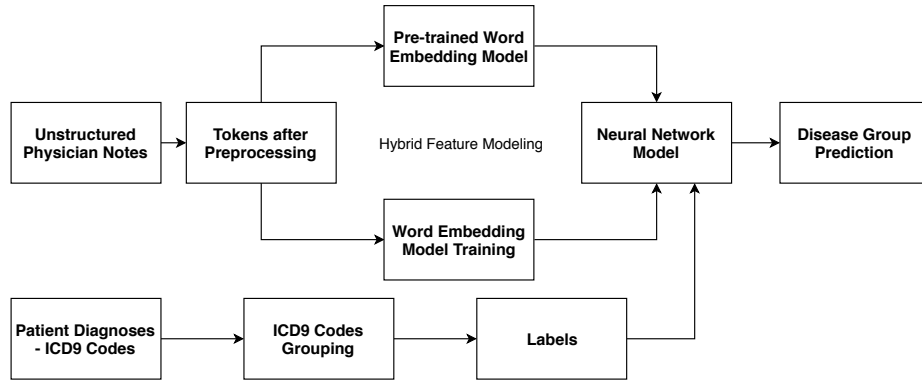


Fig. 1. Proposed ICD9 Disease Group Prediction Process

identified errors were removed from the cohort. Additionally, those records with less than 15 words are removed and finally, the remaining 141,209 records are considered for the study. Some characteristic features of the physician notes corpus are tabulated in Table 1.

We observed that there were 832 kinds of physician notes available in the MIMIC-III dataset such as Physician Resident Progress note, Intensivist note, etc. The frequency statistics of the top ten kinds of physician notes are tabulated in Table 2. A particular patient may suffer from multiple diagnoses during a particular admission and hence, it is necessary for the prediction to be a multi-label prediction task. Therefore, for each physician note, all the diagnosed disease groups during that particular admission were considered as labels and given binary values - 0 (if the disease was not diagnosed) and 1 (if the disease was diagnosed).

Table 1. Physician Notes Corpus Characteristics

Feature	Total Records
Physician Notes	141,209
Unique Words	635,531
Words in longest Note	3,443
Words in shortest Note	16
Average word length of notes	858
Unique Diseases	4,208
Disease Groups	20

**Table 2.** Top 10 Physician Notes Types

<b>Note Type</b>	<b>Occurrences</b>
Physician Resident Progress	62,550
Intensivist	26,028
Physician Attending Progress	20,997
Physician Resident Admission	10,611
ICU Note - CVI	4,481
Physician Attending Admission (MICU)	3,307
Physician Resident/Attending Progress (MICU)	1,519
Physician Surgical Admission	1,102
Physician Fellow/Attending Progress (MICU)	970
Physician Attending Admission	873

## 2.1 Preprocessing and Feature Modeling

The physician notes corpus is first preprocessed using basic Natural Language Processing (NLP) techniques such as tokenization and stop word removal. Using tokenization, the clinical text corpus is broken down into basic units called tokens and by the stopping process, unimportant words are filtered out. The preprocessed tokens are then fed into a pre-trained word embedding model to generate the word embedding vector representation of the corpus that can be considered as textual features. The pre-trained word embedding model used in this study is an openly available model that is trained on biomedical articles in PubMed and PMC along with texts extracted from an English Wikipedia dump [20], hence capturing relevant terms and concepts in the biomedical domain, which helps generate quality feature representation of the underlying corpus. The pre-trained model generates word embedding vectors of size 1x200 for each word and these vectors were averaged to generate a representation such that each preprocessed physician note is represented as a 1x200 vector. The preprocessed tokens are then used to train a Word2Vec model [18], a neural network based word representation model that generates word embeddings based on co-occurrence of words. The Skipgram model of Word2Vec was used for training the physician notes tokens with a dimension size of 200 (same as the pre-trained model) with an initial learning rate of 0.025. The averaged word vector representation for each report tokens are extracted and then fed into the neural network model along with the vector representation extracted from the pre-trained model and the ICD9 disease group labels.

## 2.2 ICD9 Disease Code Grouping

ICU Patients' diagnoses in terms of ICD9 disease codes from the 'DIAGNOSES\_ICD' table of MIMIC-III dataset were grouped as per standards<sup>2</sup> (similar to the approach adopted by Purushotham et al. [21]). A total of 4,208 unique ICD9 disease codes thus obtained were grouped into 20 ICD9 disease groups, i.e., potential labels. As the ICD9 group prediction task is considered as a binary classification of multiple labels, 20 labels (disease groups) were considered with binary values: 0 (negative diagnosis of the disease) and 1 (positive diagnosis of the disease). The physician notes modeled into two feature matrices of shape 141209 x 200 each, along with 20 ICD9 disease groups (labels) is now used to train the neural network model.

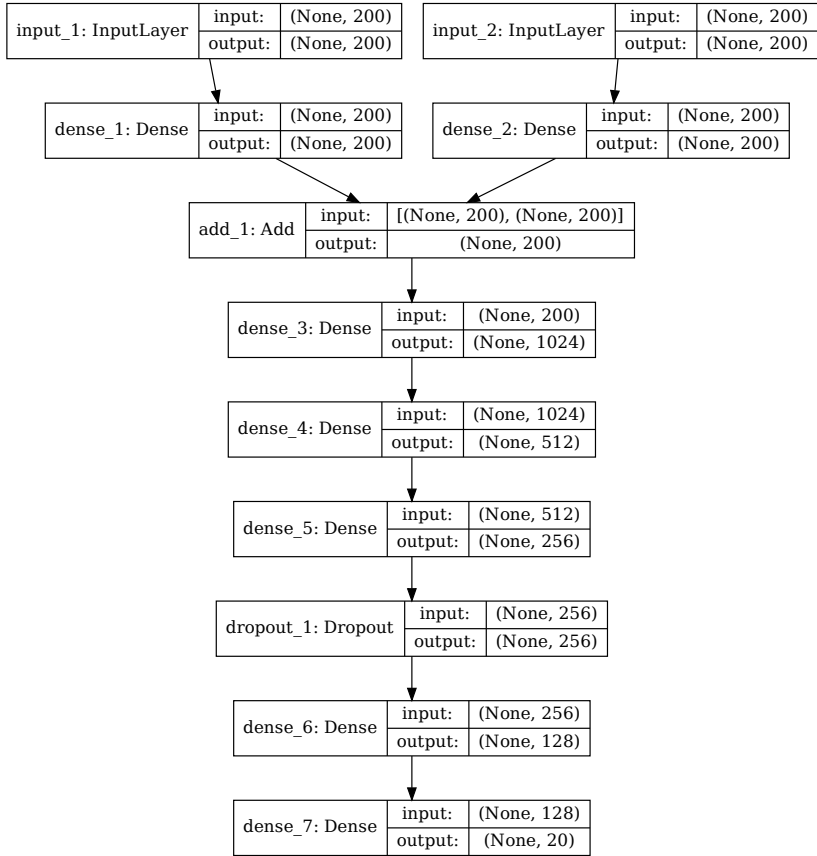
## 2.3 Neural Network Model

The proposed Deep Neural Network Prediction Model is illustrated in Fig. 2. The neural network architecture is divided into 2 parts – the first for determining the weights for the hybrid combination of features dynamically and the next for multi-label classification of ICD9 disease codes. The process of dynamically modeling the weightage to be assigned for the combination of pre-trained word embeddings and the word embeddings generated using the physicians notes is performed as shown in Fig. 3. The two feature set are fed as inputs into the neural network model, where both the input layers consists of 200 neurons, equal to the number of features generated from both models. The addition layer merges the two sets of input features in a weighted combination which is dynamically determined through backpropagation of the overall neural network architecture thereby ensuring the optimal and effective combination that offers the best classification performance possible. This architecture also ensures that the weights for the hybrid combination of features is always determined dynamically and hence can be used for any clinical text corpus. The combined set of features, i.e., the hybrid features, are then fed on to the dense feed forward neural network model which performs the training for multi-label classification of ICD9 disease code groups.

## 2.4 Disease Prediction Model

The dynamically weighted feature matrix consisting of the hybrid word embedding features, along with ICD9 group labels are next used for training a Feed Forward Neural Network (FFNN) used as the prediction model (depicted in Fig. 4). The input layer consists of 1024 neurons with input dimension as 200 (number of input features); followed by three hidden layers with 512, 256 and 128 neurons respectively and finally an output layer with 20 neurons, each representing an ICD-9 disease group. To prevent overfitting, a dropout layer, with a dropout rate of 20% was also added to the FFNN model (see Fig. 2). As this is a

<sup>2</sup> Available online [http://tdrdata.com/ipd/ipd\\_SearchForICD9CodesAndDescriptions.aspx](http://tdrdata.com/ipd/ipd_SearchForICD9CodesAndDescriptions.aspx)



**Fig. 2.** Overall Neural Network Model for ICD9 Group Prediction

binary classification for multiple labels, binary cross entropy was used as a loss function, while Stochastic Gradient Descent (SGD) was used as the optimizer with a learning rate of 0.01. Rectified Linear Unit *ReLU* activation function was used as the input and hidden layer activation functions as the feature matrix values are standardized to the range -1 and 1. The major hyperparameters for the FFNN model – the optimizer, learning rate of the optimizer and the activation function, were tuned empirically over several experiments using the GridSearchCV function in Python sklearn library. Finally, the output layer activation function is a sigmoid, again as the classification is two-class for each of the 20 labels. Training was performed for 50 epochs and then the model was applied to the validation set to predict disease groups after which the results were observed and analyzed.

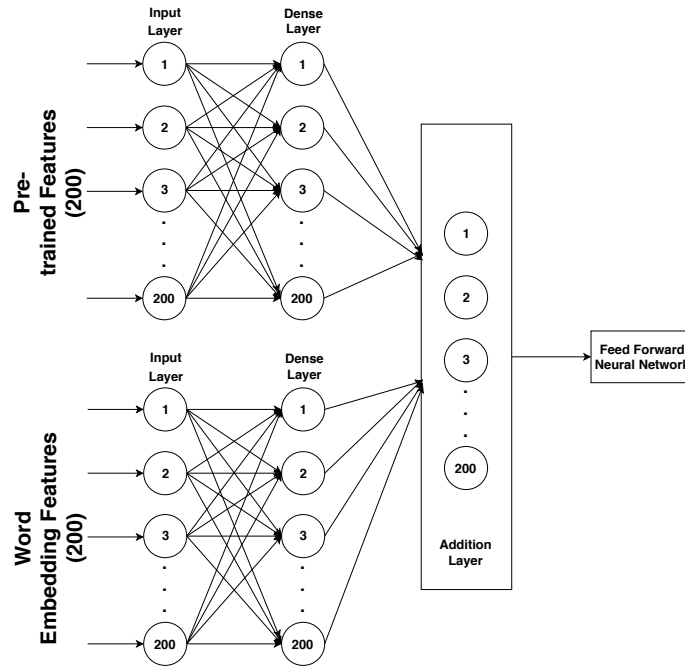


Fig. 3. Hybrid Features using Dynamic Weighted Addition of Feature Representations

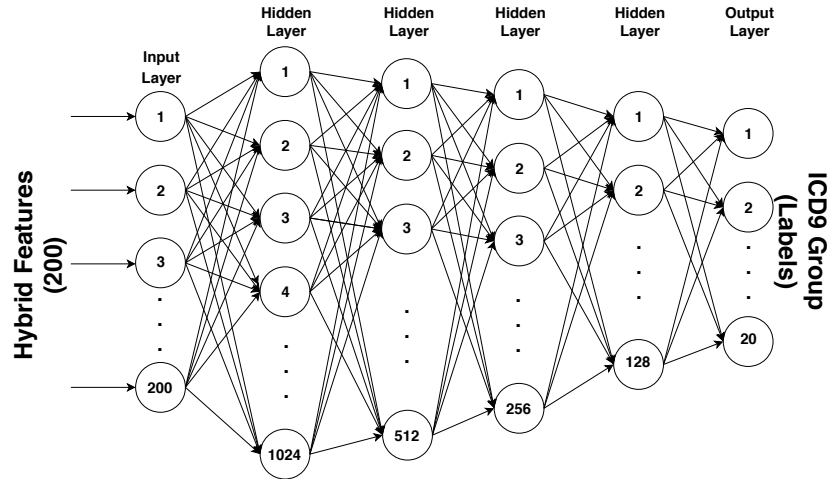


Fig. 4. Feed-forward Neural Network - Disease Group Prediction Model

### 3 Experimental Results and Discussion

To evaluate the proposed approach, we performed several experiments using standard metrics like accuracy, precision, recall, F-score, Area under Receiver



Operating Characteristic curve (AUROC), Area under Precision Recall Curve (AUPRC) and Matthew’s Correlation Coefficient (MCC). We measured these metrics on a sample-wise basis, i.e., for each report, the predicted and actual disease groups were compared and analyzed. It can be observed from the Table 3 that, the proposed model achieved promising results: AUPRC of 0.85 and AUROC of 0.89. The accuracy, precision, MCC and F-score of 0.79, 0.82, 0.57 and 0.79 respectively also indicate a good performance.

For evaluating the proposed hybrid text feature modeling approach, we designed experiments to compare its performance against that of baseline feature modeling approaches – TF-IDF based bag-of-words approach, trained word embedding approach (only Word2Vec model) and a pre-trained word embedding approach (using word embedding model trained on PubMed, PMC and Wikipedia English articles). The bag-of-words approach is based on calculated term Frequency and inverse document frequency (*tfXidf*) scores determined from the physician notes corpus. The Sklearn English stopwords list was used to filter the stopwords and  $n$ -gram ( $n = 1, 2, 3$ ) features were considered. Finally, the top 1000 features were extracted from the corpus and then fed into the neural network model for training. The other two baselines were kept the same as explained in Section 2. It is to be noted that in the neural network configuration, only one input is present, i.e., there is no hybrid weighted addition layer. The results of comparison are tabulated in Table 3. It can be observed that the proposed approach that involves a hybrid weighted combination of pre-trained and trained word embeddings is able to perform comparatively better in terms of all metrics.

Next, a comparative benchmarking against the state-of-the-art ICD9 disease group prediction model developed by Purushotham et al. [21] was performed. For each hospital admission ID (HADM\_ID) considered by Purushottam et al [21], the physician notes (if present) were extracted from ‘NOTEVENTS’ table. Although the number of records under consideration for both the studies are different, it is to be noted that the labels are distributed similarly (statistics shown in Fig. 5) and is therefore considered a fair comparison. We consider this comparison in order to study the effect of the disease group prediction models built on structured patient data (state-of-the-art approach [21]) and unstructured patient data (physician notes in this case). The results of the benchmarking are tabulated in Table 4, which clearly shows the proposed approach outperformed Purushotham et al’s model Purushotham et al. [21] by 15% in terms of AUROC and 40% in terms of AUPRC. This shows that the predictive power of a model built on unstructured patient data exceeds that of those built on structured data, where some relevant information may be lost during the coding process. To encourage comparative studies that use physician notes in MIMIC-III dataset, certain additional metrics were also considered. The Recall & F-Score performance as well as the MCC values of the proposed model over our easily reproducible patient cohort data subset are also observed and provided.


**Table 3.** Experimental Results – Baseline Comparison

Parameter	Proposed Approach	Bag-of-words (TF-IDF)	Only Word2Vec	Only Pre-trained
AUROC	0.89	0.87	0.88	0.88
AUPRC	0.85	0.81	0.84	0.82
Accuracy	0.79	0.78	0.79	0.79
Precision	0.82	0.80	0.82	0.80
Recall	0.79	0.78	0.79	0.78
F-Score	0.79	0.78	0.79	0.79
MCC	0.58	0.53	0.57	0.56

**Table 4.** Experimental Results – Comparison with the State-of-the-art

Parameter	Our Approach	Purushotham et al. [21]
Type of Data	Unstructured text	Structured data
AUROC	0.89	0.77
AUPRC	0.85	0.60
Accuracy	0.79	*
Precision	0.82	*
Recall	0.79	*
F-Score	0.79	*
MCC	0.58	*

\* Results not reported in the study

**Fig. 5.** ICD9 Group Labels Statistics Comparison

*Discussion.* From our experiments, we observed a significant potential in developing prediction based CDSS using unstructured text reports directly, eliminating the dependency on the availability of structured patient data and EHRs. The proposed approach that involves a textual feature modeling and a neural network based prediction model was successful in capturing the rich and latent clinical information available in unstructured physician notes, and using it to effectively learn disease group characteristics for prediction. The Word2Vec model was trained on the physician notes corpus with optimized parameter configuration generates effective word embedding features to be fed into the neural network model. The hybrid combination of these with the embedding features of the corpus generated by the pre-trained model trained on PubMed and PMC articles further enhanced and enriched the semantics of the textual features with biomedical domain knowledge. This is clear from the baseline comparison shown in Table 3 and it is this combination that has further enabled the FFNN to generalize better and learn the textual feature representation, effectively improving prediction performance when compared to the state-of-the-art model built on structured data.

It is interesting to note that the patient data was modeled using only textual features, without any EHRs, structured data or other processed information. The high AUPRC and AUROC values obtained in comparison to the state-of-the-art's (based on structured data) performance is an indication that the unstructured text clinical notes (physician notes in this case) contain abundant patient-specific information that was beneficial for predictive analytics applications. Moreover, the conversion process from unstructured patient text reports to structured data can be eliminated, thereby saving huge man hours, cost and other resources. The proposed approach also eliminates any dependency on structured EHRs, thus making it suitable for deployment in developing countries.

We found that the data preparation pipeline adopted for this study could be improved, as it sometimes resulted in some conflicting cases during training. This is because of the structure of the MIMIC-III dataset, in which, the physician notes do not have a direct assignment to ICD9 disease codes. To overcome this problem, we designed an assignment method in which we extracted ICD9 codes from the DIAGNOSES\_ICD table and assigned them to all patients with the same SUBJECT\_ID and HADM\_ID in the physician notes data subset. A side-effect of this method is that, at times, the ICD9 disease codes/groups assigned to the physician notes may not be related to that particular disease as a patient under same admission can have multiple physician notes. However, the model achieved promising results and the good performance indicates that prediction model was able to capture disease-specific features using the information present in the notes during the admission.

## 4 Conclusion and Future Work

In this article, a deep neural network based model for predicting ICD9 disease groups from physician notes in the form of unstructured text is discussed. The

proposed approach is built on a hybrid feature set consisting of word embedding features generated from a Word2Vec Skipgram model and also from a pre-trained word embedding model trained on biomedical articles from PubMed and PMC. The ICD9 disease codes were categorized into 20 standard groups and then used to train a binary classifier for multi-label prediction. The two sets of word embedding features were input into a neural network model as inputs and the weights of hybrid addition was determined dynamically using backpropagation. The combined features were further fed into a FFNN architecture to train the classifier and the prediction model was validated and benchmarked against state-of-the-art ICD9 disease group prediction model. The experiments highlighted the promising results achieved by the proposed model, outperforming the state-of-the-art model by 15% in terms of AUROC and 40% in terms of AUPRC.

The promising results achieved by the proposed model underscores its usefulness as an alternative to current CDSS approaches, as it eliminates the dependency on structured clinical data, ensuring that hospitals in developing countries with low structured EHR adoption rate can also make use of effective CDSS in their functioning. As part of future work, we plan to address the issues observed in the current data preparation strategy, and enhance it by sorting out the disease group assignment problems. We also intend to explore other techniques to further optimize topic and feature modeling of textual representations to study their effect on disease prediction. Additionally, we intend to evaluate and compare other disease group prediction approaches on the patient cohort used in this study. Furthermore, we also aim to perform a detailed efficiency analysis and comparison with respect to training times of the proposed approach and various baselines. Finally, extensive benchmarking of approaches for modeling unstructured clinical notes on the physician notes, for further insights into the adaptability of the proposed approach in real-world scenarios will also be explored.

## Acknowledgement

We gratefully acknowledge the use of the facilities at the Department of Information Technology, NITK Surathkal, funded by Govt. of India's DST-SERB Early Career Research Grant (ECR/2017/001056) to the second author.

## References

- [1] Appelros, P.: Prediction of length of stay for stroke patients. *Acta Neurologica Scandinavica* 116(1), 15–19 (2007)
- [2] Ayyar, S., Don, O., Iv, W.: Tagging patient notes with icd-9 codes. In: *Proceedings of the 29th Conference on Neural Information Processing Systems* (2016)
- [3] Baumel, T., Nassour-Kassis, J., Cohen, R., Elhadad, M., Elhadad, N.: Multi-label classification of patient notes: case study on icd code assign-

- ment. In: Workshops at the Thirty-Second AAAI Conference on Artificial Intelligence (2018)
- [4] Berndorfer, S., Henriksson, A.: Automated diagnosis coding with combined text representations. *Studies in health technology and informatics* 235, 201 (2017)
- [5] Calvert, J., Mao, Q., Hoffman, J.L., Jay, M., Desautels, T., Mohamadlou, H., Chettipally, U., Das, R.: Using electronic health record collected clinical variables to predict medical intensive care unit mortality. *Annals of Medicine and Surgery* 11, 52–57 (2016)
- [6] Choi, E., Bahadori, M.T., Schuetz, A., Stewart, W.F., Sun, J.: Doctor ai: Predicting clinical events via recurrent neural networks. In: *Machine Learning for Healthcare Conference*. pp. 301–318 (2016)
- [7] Farkas, R., Szarvas, G.: Automatic construction of rule-based icd-9-cm coding systems. In: *BMC bioinformatics*. vol. 9, p. S10. BioMed Central (2008)
- [8] Gangavarapu, T., Krishnan, G., Kamath, S., Jeganathan, J.: Farsight: Long-term disease prediction using unstructured clinical nursing notes. *IEEE Transactions on Emerging Topics in Computing* (01), 1–1 (2020)
- [9] Gangavarapu, T., Jayasimha, A., Krishnan, G.S., Kamath, S.: Predicting icd-9 code groups with fuzzy similarity based supervised multi-label classification of unstructured clinical nursing notes. *Knowledge-Based Systems* 190, 105321 (2020)
- [10] Ge, W., Huh, J.W., Park, Y.R., Lee, J.H., Kim, Y.H., Turchin, A.: An interpretable icu mortality prediction model based on logistic regression and recurrent neural networks with lstm units. In: *AMIA Annual Symposium Proceedings*. vol. 2018, p. 460. American Medical Informatics Association (2018)
- [11] Harutyunyan, H., Khachatryan, H., Kale, D.C., Galstyan, A.: Multitask learning and benchmarking with clinical time series data. *arXiv preprint arXiv:1703.07771* (2017)
- [12] Jiang, S., Chin, K.S., Qu, G., Tsui, K.L.: An integrated machine learning framework for hospital readmission prediction. *Knowledge-Based Systems* 146, 73–90 (2018)
- [13] Johnson, A.E., Pollard, T.J., Shen, L., Li-wei, H.L., Feng, M., Ghassemi, M., Moody, B., Szolovits, P., Celi, L.A., Mark, R.G.: Mimic-iii, a freely accessible critical care database. *Scientific data* 3, 160035 (2016)
- [14] Kansagara, D., Englander, H., Salanitro, A., Kagen, D., Theobald, C., Freeman, M., Kripalani, S.: Risk prediction models for hospital readmission: a systematic review. *Jama* 306(15), 1688–1698 (2011)
- [15] Krishnan, G.S., Kamath S, S.: Ontology-driven text feature modeling for disease prediction using unstructured radiological notes. *Computación y Sistemas* 23(3) (2019)
- [16] Krishnan, G.S., S., S.K.: A novel ga-elm model for patient-specific mortality prediction over large-scale lab event data. *Applied Soft Computing* 80, 525 – 533 (2019), <http://www.sciencedirect.com/science/article/pii/S1568494619302108>

- [17] Li, M., Fei, Z., Zeng, M., Wu, F., Li, Y., Pan, Y., Wang, J.: Automated icd-9 coding via a deep learning approach. *IEEE/ACM transactions on computational biology and bioinformatics* (2018)
- [18] Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013)
- [19] Miotto, R., Li, L., Kidd, B.A., Dudley, J.T.: Deep patient: an unsupervised representation to predict the future of patients from the electronic health records. *Scientific reports* 6, 26094 (2016)
- [20] Nédellec, C., Bossy, R., Kim, J.D., Kim, J.J., Ohta, T., Pyysalo, S., Zweigenbaum, P.: Overview of bionlp shared task 2013. In: *Proceedings of the BioNLP Shared Task 2013 Workshop*. pp. 1–7 (2013)
- [21] Purushotham, S., Meng, C., Che, Z., Liu, Y.: Benchmarking deep learning models on large healthcare datasets. *Journal of biomedical informatics* (2018)
- [22] Reddy, B.K., Delen, D.: Predicting hospital readmission for lupus patients: An rnn-lstm-based deep-learning methodology. *Computers in biology and medicine* 101, 199–209 (2018)
- [23] Shickel, B., Loftus, T.J., Adhikari, L., Ozrazgat-Baslanti, T., Bihorac, A., Rashidi, P.: Deepsofa: A continuous acuity score for critically ill patients using clinically interpretable deep learning. *Scientific reports* 9(1), 1879 (2019)
- [24] Van Houdenhoven, M., Nguyen, D.T., Eijkemans, M.J., Steyerberg, E.W., Tilanus, H.W., Gommers, D., Wullink, G., Bakker, J., Kazemier, G.: Optimizing intensive care capacity using individual length-of-stay prediction models. *Critical Care* 11(2), R42 (2007)
- [25] Xie, P., Xing, E.: A neural architecture for automated icd coding. In: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. vol. 1, pp. 1066–1076 (2018)
- [26] Zeng, M., Li, M., Fei, Z., Yu, Y., Pan, Y., Wang, J.: Automatic icd-9 coding via deep transfer learning. *Neurocomputing* 324, 43–50 (2019)