Multimodal Learning for Cardiovascular Risk Prediction using EHR Data

Ayoub Bagheri*
a.bagheri@uu.nl
Methodology and Statistics
Utrecht University
Utrecht, The Netherlands

T. Katrien J. Groenhof
t.k.j.groenhof@umcutrecht.nl
Julius Center for Health Sciences and
Primary Care
University Medical Center Utrecht
Utrecht, The Netherlands

Wouter B. Veldhuis
w.veldhuis@umcutrecht.nl
Department of Radiology
University Medical Center Utrecht
Utrecht, The Netherlands

Pim A. de Jong
pjong8@umcutrecht.nl
Department of Radiology
University Medical Center Utrecht
Utrecht, The Netherlands

Folkert W. Asselbergs[†‡]
f.w.asselbergs@umcutrecht.nl
Department of Cardiology
University Medical Center Utrecht
Utrecht, The Netherlands

Daniel L. Oberski[§]
d.l.oberski@uu.nl
Methodology and Statistics
Utrecht University
Utrecht, The Netherlands

Abstract

Electronic health records (EHRs) contain structured and unstructured data of significant clinical and research value. Various machine learning approaches have been developed to employ information in EHRs for risk prediction. The majority of these attempts, however, focus on structured EHR fields and lose the vast amount of information in the unstructured texts. To exploit the potential information captured in EHRs, in this study we propose a multimodal recurrent neural network model for cardiovascular risk prediction that integrates both medical texts and structured clinical information. The proposed multimodal bidirectional long short-term memory (BiLSTM) model concatenates word embeddings to classical clinical predictors before applying them to a final fully connected neural network. In the experiments, we compare performance of different deep neural network (DNN) architectures including convolutional neural network and long short-term memory in scenarios of using clinical variables and chest X-ray radiology reports. Evaluated on a data set of real world patients with manifest vascular disease or at high–risk for cardiovascular disease, the proposed BiLSTM model demonstrates state-of-the-art performance and outperforms other DNN baseline architectures.

Keywords: clinical text mining, text classification, cardiovascular risk prediction, clinical multimodal learning

1 Introduction

Electronic health records (EHRs) data have become increasingly available to researchers as more hospitals, clinics and practices have adopted data digitization. EHRs store data in different modalities, such as structured tabular data (e.g. demographic values, laboratory results, medications) and unstructured texts (e.g. referral letters, clinical notes, discharge summaries, radiology reports). This digitization creates an opportunity to mine the health records to increase the quality of care and clinical outcomes. Yet clinicians have limited time to process all the available data and detect patterns across similar medical records. Deep learning and machine learning, on the other hand, are suitable for discovering useful patterns from vast amount of data.

Unstructured texts contained within the EHRs are recognized as a rich but not easily accessible and usable source of medical information [10, 34, 42, 47]. Recent studies have attempted to derive information from unstructured medical texts to classify disease codes [3, 11], detect patient's disease history [2, 44], and predict hospital readmission or clinical outcomes [1, 16, 17]. X-ray radiology reports are example of such unstructured data describing radiologist's observations on patient's medical conditions associated to medical images. The majority of previous decision support systems for radiology reports are developed using rule-based approaches applied on unstructured and semi-structured texts [8, 13, 39, 41]. However, these methods are often impractical because they do not generalize to new data and often are not applicable for big data analysis [28].

Recent studies have shown promising results using free-text radiology reports and deep learning models to predict clinical outcomes [8, 21, 26, 35, 37, 43]. Convolutional neural networks (CNNs) and recurrent neural networks (RNNs) are two common deep learning techniques that have been effective in text mining and natural language processing (NLP), also in EHR applications [4, 11, 17, 26]. Deep learning-based modelling of radiology reports has been proposed to supersede the simple grammatical patterns and hand-crafted regular expressions of the traditional clinical rules-based

---

*Also with Department of Cardiology, University Medical Center Utrecht.
†Also with Institute of Cardiovascular Science, University College London.
‡Also with Health Data Research UK and Institute of Health Informatics, University College London.
§Also with Julius Center for Health Sciences and Primary Care, University Medical Center Utrecht.

software, such as PEFinder [7], medical language extraction and encoding system (MedLEE) [12, 15] and CTakes [30]. While these neural networks models gained tremendous momentum in knowledge discovery from EHR texts, there are very seldom studies that used of both free-texts and the structured information in EHRs for clinical prediction and classification [18, 23, 31, 46].

In this paper we leverage structured features in EHR data, i.e. lab results, to combine with free-text radiology reports to uncover patterns to improve cardiovascular risk prediction. The primary contribution of this study is twofold:

- We present a multimodal bi-directional long short-term memory (BiLSTM) neural network that integrates the neural text representation with laboratory results and feeds them into a fully connected neural network.
- We investigate the effectiveness of the proposed architecture to predict cardiovascular risk for real world patients with manifest vascular disease or at high-risk for cardiovascular disease [36] visiting the University Medical Center (UMC) Utrecht.

The rest of paper is organized as follows, section 2 provides related work in the literature, section 3 details the proposed multimodal architecture for mining EHR data. Experiments and results are given in section 4. Section 5 concludes this research.

2   Related Work

Many text mining pipelines and NLP systems have been developed to extract structured information from free-text radiology reports. These range from simple rule-based, regular expression, and bag-of-word methods, to more sophisticated machine learning- and deep learning-based approaches [4, 5, 8–10, 12, 13, 21, 26, 28, 33, 35, 37, 39, 43, 45, 47].

Sevenster et al. [33] evaluated a system that extracts and correlates clinical findings and body locations from radiology reports using MedLEE, a rule-based NLP-based system. Khalifa and Meystre [19] built a NLP application based on the Apache UIMA (Unstructured Information Management Architecture) and reusing existing tools previously developed. Using this application they addressed identifying risk factors for heart disease based on the automated analysis of narrative clinical records of diabetic patients. Drozdov et al. [10] applied five supervised machine learning algorithms on chest radiology reports; K-nearest neighbour, logistic regression, Gaussian naÃŕve Bayes classifier, random forest, and support vector machine. These methods were evaluated on a term frequency-inverse document frequency matrix of radiology reports that was then transformed to lowered dimensions using singular value decomposition. These rule-based and traditional machine learning approaches are in the need of domain expertise and hard core feature extraction [28, 34, 41].

Deep learning techniques have already shown potential to automate the task of classifying X-ray reports in a way that could inform decisions regarding medical utilization [8, 10, 35]. Recently, a supervised learning approach using a RNN model with attention mechanism achieved high accuracy on expert-labeled chest X-ray radiology reports data set [5]. Similarly, Cornegruta et al. [9] proposed a BiLSTM neural network which was demonstrated to perform favourably in a corpus of radiology reports. Drozdov et al. [10], besides non-neural classifiers, applied LSTM and BiLSTM networks on chest X-ray radiology reports that could produce state-of-the-art classification results. Wood et al. [43] developed ALARM, a transformer-based network report classifier on MRI data using BioBERT [22] models trained on radiology reports, and demonstrated improvement over simpler word embedding methods [25, 27, 47]. Smit et al. [37] demonstrated superior performance of a deep learning method for radiology report labeling, in which a biomedically pretrained BERT model is first trained on annotations of a rule-based labeler, and then finetuned on a small set of expert annotations. Finally, Chen et al. [8] implemented CNNs to extract pulmonary embolism findings from thoracic computed tomography reports, outperforming state-of-the-art NLP systems.

However, current studies on text mining for radiology reports lack the flexibility to utilize the benefits of using multimodal data, i.e. both unstructured text and structured numerical values. One study using both laboratory features and text reports is that of Liu et al. [23]. Liu et al. concatenated lab and demographic features to the output features of medical notes obtained from a max-pooling layer of their neural network model. In their study, CNN- and RNN-based deep learning models have been developed for chronic disease prediction using medical notes, and then a fully connected neural network has been applied with one hidden layer at the end of the pipeline. Another similar research is by Xu et. al [46], in which separate machine learning models were trained with data from unstructured text, semi-structured text and structured tabular data to create a model that predicts diagnostic codes of international classification of diseases (ICD-10).

In this paper, we present a multimodal learning architecture with a bidirectional deep learning-based model for free-text radiology reports. The bidirectional model can deal with contextual dependencies in text as it trains two LSTM networks on each side of the current word, running from left to right and from right to left, to encode and represent the text. The research presented in this paper is different from previous research as it: (1) combines structured laboratory results and unstructured text in a machine learning model to make an accurate decision of classifying cardiovascular events. (2) uses real world EHR data including X-ray reports in Dutch for cardiovascular risk prediction.

3 Methodology

In this section, we describe the details of our proposed methods, including the used cohort study, data ethics and privacy, text preprocessing, and the employed multimodal recurrent neural network.

3.1 Cohort Study

The patients included in this study originated from the Second Manifestations of ARTerial disease (SMART) study. The design of the SMART study is published elsewhere [36]. In short, the SMART study is an ongoing single-center prospective cohort study that was designed to establish the presence of additional arterial disease and risk factors for atherosclerosis in patients with manifest vascular disease or a vascular risk factor. Patients visiting the UMC Utrecht for evaluation of an atherosclerotic cardiovascular condition are eligible for inclusion in SMART. Inclusion criteria are presentation with an atherosclerotic cardiovascular condition and age > 18 years. Exclusion criteria are life expectancy < 3 months, unstable vascular disease and insufficient fluency in the Dutch language. A total of 5603 patients and their reports were included in this analysis. Characteristics are shown in Table 1. Variables that are predictors in the SMART risk score (age, sex, smoking, systolic blood pressure, diabetes, high-density-lipoprotein (HDL) cholesterol, total cholesterol, renal function according to the modification of diet in renal disease (MDRD) formula, history of cardiovascular disease stratified for stroke, peripheral artery disease, abdominal aortic aneurysm, and coronary heart disease, and years since diagnosis of first cardiovascular disease) were extracted for all patients. MACE (MAjor Cardiovascular Events) during followup is the outcome variable for prediction. Missingness of data was solved using the MICE package [6] with one imputation for each missing value.

3.2 Ethics and Privacy

Informed consent was obtained through established procedures. The SMART study has been approved by the Medical Ethical Committee of the UMC Utrecht. All data are handled according to data protection and privacy regulations.

3.3 Text Preprocessing

Radiology reports in the SMART study, for the process of text mining, may contain redundant characters and words such as punctuation marks and stop words. Therefore, we perform the following preprocessing steps to improve the quality of text data for the ongoing study: (1) All characters are transformed into lowercase. (2) We remove numbers and some meaningless punctuation marks such as semicolon and colon. (3) Stop words are then removed. (4) We then apply the Porter's stemming algorithm [20, 29] on texts.

Table 1. Patients' Characteristics in the SMART Study

Characteristic	Total n = 5603
Age, years (mean (sd))	56.2 (12.5)
Female sex, n (%)	1926 (34.4)
Current smoking, n (%)	1549 (27.6)
History of cardiovascular disease	
CHD*, n (%)	2166 (38.66)
Stroke, n (%)	1076 (19.20)
PAD*, n (%)	631 (11.26)
AAA*, n (%)	306 (5.46)
Years since first diagnosis of CVD*, median (IQR)	0 (0-4)
Risk factors for cardiovascular disease	
Diabetes Mellitus, n (%)	1047 (18.69)
Hypertension, n (%)	2353 (42.00)
Dyslipidemia, n (%)	432 (7.71)
BMI*, kg/m2 (mean (sd))	26.8 (4.3)
SBP*, mmHg (mean (sd))	140 (21)
DBP*, mmHg (mean (sd))	83 (13)
Laboratory	
Total cholesterol, mmol/L (mean (sd))	5.14 (1.38)
LDL*-cholesterol, mmol/L (mean (sd))	3.1 (1.16)
HDL-cholesterol, mmol/L (mean (sd)	1.27 (0.38)
Triglycerides, mmol/L (median (IQR))	1.7 (1.2-2.5)
MDRD, ml/min/1.73m2 (median (IQR))	80 (68-91)
HbA1c*, mmol/mol (median (IQR))	5.7 (5.4-6.1)
Glucose, mmol/L (median (IQR))	5.7 (2.6-6.4)
Hemoglobin, mmol/L (mean (sd))	6.0 (2.04)
Creatinine, μmol/L (median (IQR))	84 (73-97)
CRP*, mg/L (median (IQR))	1.95 (0.90-4.20)
TSH*, mU/l (mean (sd))	0.9 (0.09)
MACE during followup, n (%)	1385 (24.72)

* CVD: Cardiovascular disease, CHD: Coronary heart disease, PAD: Peripheral arterial disease, AAA: Abdominal aortic aneurysm, BMI: Body mass index, SBP: Systolic blood pressure, DBP: Diastolic blood pressure, LDL: Low-density lipoprotein, HbA1c: Hemoglobine A1c, CRP: C-reactive protein, TSH: Thyroid stimulating hormone

3.4 Multimodal Recurrent Neural Network

Our proposed model consists of an embedding layer, a BiLSTM layer, dropout, a concatenation layer and dense layers. Figure 1 shows the architecture of the multimodal BiLSTM model that integrates word embeddings and clinical predictors.

3.4.1 Embedding Layer. To extract the semantic information of radiology reports, each text is firstly represented as a sequence of word embeddings. Word embedding is an improvement over bag-of-word models where large sparse vectors were used to represent each word. On the contrary, in an embedding, words are represented by dense vectors
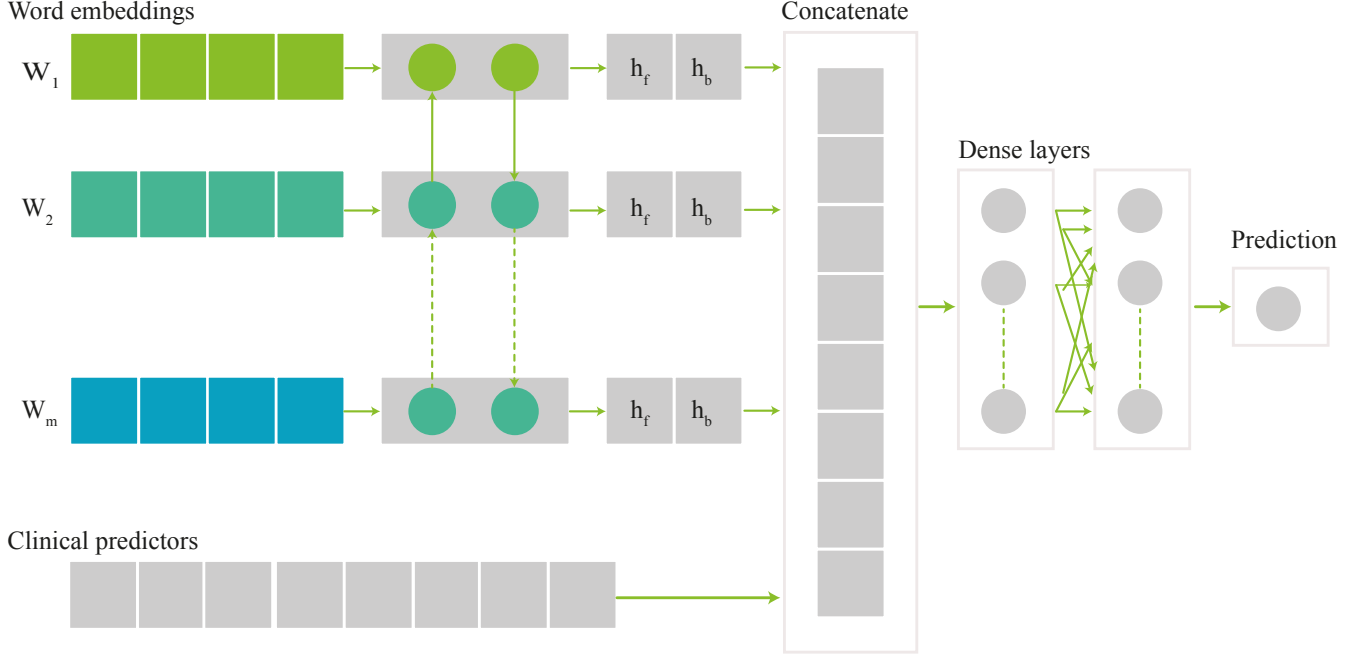
Figure 1. Multimodal Recurrent Neural Network

where a vector represents the projection of the word into a continuous vector space [24, 25]. Denote s as an X-ray report with m words and each word is mapping to a vector, then we have:

$$s = [\vec{e}_1, \vec{e}_2, ..., \vec{e}_m] \tag{1}$$

where vector $\vec{e}_i$ represents the vector of i-th word with a dimension of d. The vectors of word embeddings are concatenated together to maintain the order of words in a patient report.

3.4.2 Bidirectional-LSTM Layer. After the embedding layer, the sequence of word vectors is fed into a bidirectional LSTM layer to achieve another representation of radiology reports. Interest in incorporating a BiLSTM layer into the architecture of our model arises from their ability to learn long-term dependencies and contextual features from both past and future states [32]. The BiLSTM layer calculates two parallel LSTM layers, a forward hidden layer and a backward hidden layer, to generate an output sequence y as illustrated:

$$h_{f_t} = \sigma(W_{xh_f} x_t + W_{h_f h_f} h_{f_{t-1}} + b_{h_f}) \tag{2}$$

$$h_{b_t} = \sigma(W_{xh_b} x_t + W_{h_b h_b} h_{b_{t-1}} + b_{h_b}) \tag{3}$$

$$y_t = W_{h_f y} h_{f_t} + W_{h_b y} h_{b_t} + b_y \tag{4}$$

Here σ is the sigmoid activation function, x_t is a d-dimensional input vector at time step t, W are the weight matrices, b are bias vectors, and h_f, h_b are the output of the LSTM forward and backward layers, respectively.

3.4.3 Dropout. Large neural networks trained on relatively small data sets can overfit the training data. Dropout provides a powerful method of regularizing that prevents the potential overfitting issue by randomly setting input elements to zero with a given probability of dropout rate [14, 38]. In our multimodal RNN model, dropout and recurrent dropout are used with the BiLSTM layer.

3.4.4 Concatenation Layer. The concatenation layer, in the multimodal RNN model, takes the outputs of the BiLSTM layer and the clinical predictors from the SMART study as inputs and concatenates them along a specified dimension. In this layer there are no weights to be learned.

3.4.5 Dense Layers. Dense layers after the concatenation layer add predictive value as they are able to learn interactions between the text features and the clinical predictors. In our multimodal model, the concatenation layer is followed by dense layers with the same sizes of neurons. The output of the concatenation layer is fed into the first dense layer with the *ReLU* activation function. The output of the second dense layer is then fed into the third dense layer with one neuron. The activation function in this layer is *sigmoid* with the *binary crossentropy* loss function.

Table 2. Hyperparameter Values

Parameter	Value
Embedding dimension (d)	500
#neurons in LSTM layer	100
CNN filter size	5
#filters in CNN	128
#neurons in dense layers	64
Dropout rate	0.2
Recurrent dropout rate	0.2
Batch size	64
#epochs	20
Optimization method	ADAM

4 Experiments

4.1 Evaluation Metrics

To evaluate the classification performance, we use five available metrics: AUC (area under the curve), misclassification rate, precision, recall and F1 score. AUC is the area under the receiver operating characteristic curve which is created by plotting the true positive rate against the false positive rate. Misclassification rate is the proportion of incorrectly classified instances made by a model. Precision is the fraction of relevant instances among the retrieved instances, while recall is the fraction of relevant instances that have been retrieved over the total amount of relevant instances. The F1 score can be interpreted as a weighted average of the precision and recall. The relative contributions of precision and recall to the F1 score are equal. The formula of precision, recall and F1 score can be defined as Eq. 5 ,6, 7:

$$\text{Precision} = \frac{\text{True Positive}}{\text{True Positive} + \text{False Positive}} \quad (5)$$

$$\text{Recall} = \frac{\text{True Positive}}{\text{True Positive} + \text{False Negative}} \quad (6)$$

$$\text{F1 score} = \frac{2 * \text{Precision} * \text{Recall}}{\text{Precision} + \text{Recall}} \quad (7)$$

4.2 Experiment Results

We compare the results of our multimodal BiLSTM model (MI-BiLSTM) to models in four experimental scenarios:

- V-NN: A neural network with three dense layers trained on variables from clinical predictors
- T-BiLSTM: A BiLSTM model trained on X-ray radiology reports
- MI-CNN: The multimodal model using one CNN layer
- MI-LSTM: The multimodal model using one LSTM layer

Our system is implemented on Keras with a TensorFlow backend [1]. We performed 5-fold cross validation for all experimental scenarios. The hyperparameter settings for training the models are shown in Table 2. These hyperparameters are tuned based on the validation set.

Table 3 presents AUC and misclassification rate of various models on the radiology data set of the UMC Utrecht SMART study. V-NN, which is a neural network with two hidden layers, is trained only on clinical variables. This scenario included age, sex, smoking, systolic blood pressure, diabetes, high-density-lipoprotein cholesterol, total cholesterol, renal function according to the MDRD formula, history of cardiovascular disease stratified for stroke, peripheral artery disease, abdominal aortic aneurysm, and coronary heart disease as predictors and MACE during followup as the outcome in prediction models. T-BiLSTM is trained only on X-ray reports; comparing to V-NN has a lower AUC and a higher misclassification rate. MI-CNN used the framework of the multimodal RNN with a one dimensional convolution layer followed by a max pooling layer, instead of a BiLSTM layer. Table 3 shows the effectiveness of combining both structured and unstructured data in this scenario. MI-LSTM is the multimodal RNN model with one LSTM layer. By comparing MI-LSTM and MI-BiLSTM, it can be seen that using a BiLSTM layer gains the highest metric values. This is because of the ability of the BiLSTM layer as it leverages the knowledge present on both sides of the current word to encode the text.

Figure 2 [2] shows the performance of the models using precision, recall and F1 score evaluation metrics. These results are the evidence of the performance of text mining techniques to extract the knowledge in radiology reports and combine them with classical clinical predictors. It is worth mentioning that the recall value of using only clinical variables in V-NN (79.4%) is higher than the corresponding values for T-BiLSTM (72%), MI-CNN (73.5%), and the MI-LSTM model (78.8%). The F1 score of the V-NN model is also admissible with the value of 77.2%. This shows the efficiency and relatedness of the laboratory results in predicting the cardiovascular risk.

As can be seen in Figure 2, the RNN models – MI-LSTM, MI-BiLSTM – that use both text reports and clinical variables are gained better F1 score values comparing to the CNN-based model and also to the scenarios in which either text reports or clinical variable has been employed. The MI-LSTM and MI-BiLSTM models have gained F1 score of 78.9% and 83.8%, respectively. Although CNNs have proven effective for deriving features from sequential data [11], RNNs are specialized for such data and can scale to much longer sequences than would be practical for neural networks without sequence−based specialization [14].

---

[1] https://keras.io
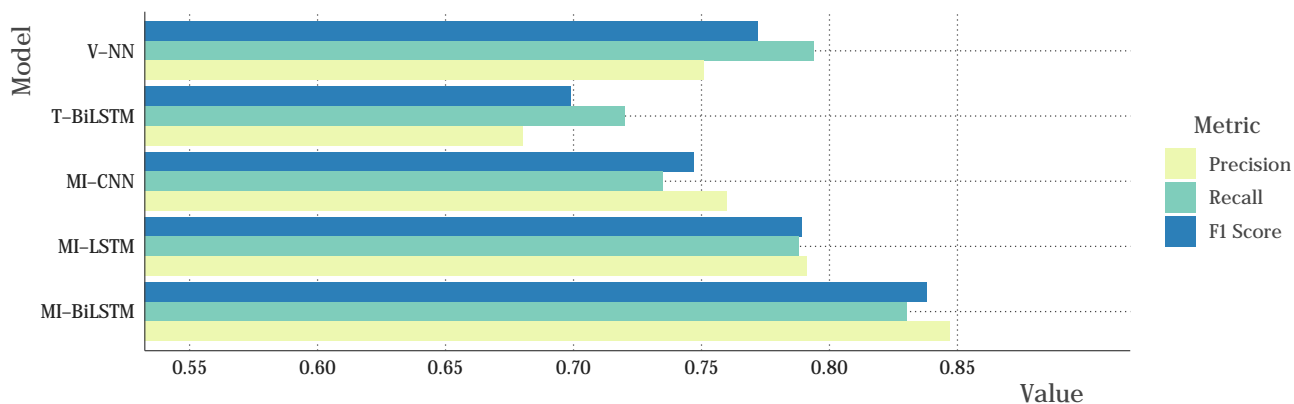[2] Firatheme version 0.2.1 [40] is used in this plot.

Figure 2. Performance Comparison of Precision, Recall and F1 Score of the Models

Table 3. AUC and Misclassification Rate of the Models on Data from the UMC Utrecht SMART Study

Model	AUC	Misclassification rate
V-NN	0.651	0.201
T-BiLSTM	0.570	0.300
MI-CNN	0.730	0.214
MI-LSTM	0.794	0.176
MI-BiLSTM	0.847	0.143

5 Conclusions

Text mining methods are the key to successful extraction of clinically important findings from radiology reports. However, previous studies for risk prediction mainly focused on modelling with either structured data or unstructured texts. In this study, we leveraged the knowledge present on both radiology reports and structured laboratory variables to predict cardiovascular risk for patients in UMC Utrecht. We proposed a neural network-driven modelling of radiological language, to integrate clinical and textual features, to supersede traditional algorithms using only clinical variables. To the best of our knowledge, there is no existing work that applies a multimodal RNN-based model to mine free-text radiology reports and combine them with laboratory results for cardiovascular risk prediction. Comparing five different scenarios, the BiLSTM RNN model showed the best performances in cardiovascular risk prediction.

Despite the great potential and the promising results of the deep learning models, the use of these advanced text mining techniques in clinical practice requires support for implementation. Implementation includes the application of the text mining pipeline, and integration in health care process using decision support systems. To help clinicians to interpret the results that come from text mining, collaborations between data scientists, software engineers and clinicians are needed to safeguard technical quality and medical relevance. To this end, in future work we plan to experiment various representations of text data in the text mining framework, and investigate the use of machine learning models with radiology data (reports and demographics) when laboratory results of patients are missing.

The publicly-available source code [3] of our model can be used to evaluate performance on clinically-relevant classification tasks based on clinical notes and EHR variables.

Acknowledgments

References

[1] Beatrice Alex, Claire Grover, Richard Tobin, Cathie Sudlow, Grant Mair, and William Whiteley. 2019. Text mining brain imaging reports. *Journal of biomedical semantics* 10, 1 (2019), 23–34.

[2] Ayoub Bagheri, Arjan Sammani, Peter Van Der Heijden, Folkert W Asselbergs, and Daniel L Oberski. 2020. ETM: Enrichment by topic modeling for automated clinical sentence classification to detect patients' disease history. *Journal of Intelligent Information Systems* (2020).

[3] Ayoub Bagheri, Arjan Sammani, Peter G.M. Van der Heijden, Folkert W. Asselbergs, and Daniel L. Oberski. 2020. Automatic ICD-10 classification of diseases from Dutch discharge letters. In *13th International Joint Conference on Biomedical Engineering Systems and Technologies*. 281–289.

[4] Imon Banerjee, Yuan Ling, Matthew C Chen, Sadid A Hasan, Curtis P Langlotz, Nathaniel Moradzadeh, Brian Chapman, Timothy Amrhein, David Mong, Daniel L Rubin, et al. 2019. Comparative effectiveness of convolutional neural network (CNN) and recurrent neural network (RNN) architectures for radiology text report classification. *Artificial intelligence in medicine* 97 (2019), 79–88.

---

[3]https://github.com/bagheria/CardioRisk-TextMining

[5] Aurelia Bustos, Antonio Pertusa, Jose-Maria Salinas, and Maria de la Iglesia-Vayá. 2019. Padchest: A large chest X-ray image dataset with multi-label annotated reports. *arXiv preprint arXiv:1901.07441* (2019).

[6] S van Buuren and Karin Groothuis-Oudshoorn. 2010. MICE: Multivariate imputation by chained equations in R. *Journal of statistical software* (2010), 1–68.

[7] Brian E Chapman, Sean Lee, Hyunseok Peter Kang, and Wendy W Chapman. 2011. Document-level classification of CT pulmonary angiography reports based on an extension of the ConText algorithm. *Journal of biomedical informatics* 44, 5 (2011), 728–737.

[8] Matthew C Chen, Robyn L Ball, Lingyao Yang, Nathaniel Moradzadeh, Brian E Chapman, David B Larson, Curtis P Langlotz, Timothy J Amrhein, and Matthew P Lungren. 2018. Deep learning to classify radiology free-text reports. *Radiology* 286, 3 (2018), 845–852.

[9] Savelie Cornegruta, Robert Bakewell, Samuel Withey, and Giovanni Montana. 2016. Modelling radiological language with bidirectional long short-term memory networks. *arXiv preprint arXiv:1609.08409* (2016).

[10] Ignat Drozdov, Daniel Forbes, Benjamin Szubert, Mark Hall, Chris Carlin, and David J Lowe. 2020. Supervised and unsupervised language modelling in chest X-ray radiological reports. *Plos one* 15, 3 (2020), e0229963.

[11] Jingcheng Du, Qingyu Chen, Yifan Peng, Yang Xiang, Cui Tao, and Zhiyong Lu. 2019. ML-Net: Multi-label classification of biomedical texts with deep neural networks. *Journal of the American Medical Informatics Association* 26, 11 (2019), 1279–1285.

[12] Carol Friedman, Philip O Alderson, John HM Austin, James J Cimino, and Stephen B Johnson. 1994. A general natural-language text processor for clinical radiology. *Journal of the American Medical Informatics Association* 1, 2 (1994), 161–174.

[13] Tianxia Gong, Chew Lim Tan, Tze Yun Leong, Cheng Kiang Lee, Boon Chuan Pang, CC Tchoyoson Lim, Qi Tian, Suisheng Tang, and Zhuo Zhang. 2008. Text mining in radiology reports. In *2008 Eighth IEEE International Conference on Data Mining*. IEEE, 815–820.

[14] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep learning*. MIT press.

[15] George Hripcsak, John HM Austin, Philip O Alderson, and Carol Friedman. 2002. Use of natural language processing to translate clinical information from a database of 889,921 chest radiographic reports. *Radiology* 224, 1 (2002), 157–163.

[16] Kexin Huang, Jaan Altosaar, and Rajesh Ranganath. 2019. ClinicalBERT: Modeling clinical notes and predicting hospital readmission. *arXiv preprint arXiv:1904.05342* (2019).

[17] Abhyuday N Jagannatha and Hong Yu. 2016. Bidirectional RNN for medical event detection in electronic health records. In *Proceedings of the conference. Association for Computational Linguistics. North American Chapter. Meeting*, Vol. 2016. NIH Public Access, 473.

[18] Mengqi Jin, Mohammad Taha Bahadori, Aaron Colak, Parminder Bhatia, Busra Celikkaya, Ram Bhakta, Selvan Senthivel, Mohammed Khalilia, Daniel Navarro, Borui Zhang, et al. 2018. Improving hospital mortality prediction with medical named entities and multimodal learning. *arXiv preprint arXiv:1811.12276* (2018).

[19] Abdulrahman Khalifa and Stéphane Meystre. 2015. Adapting existing natural language processing resources for cardiovascular risk factors identification in clinical notes. *Journal of biomedical informatics* 58 (2015), S128–S132.

[20] Wessel Kraaij and Renée Pohlmann. 1994. PorterâĂŹs stemming algorithm for Dutch. *Informatiewetenschap 1994: Wetenschappelijke bijdragen aan de derde STINFON Conferentie* (1994), 167–180.

[21] Jonathan Laserson, Christine Dan Lantsman, Michal Cohen-Sfady, Itamar Tamir, Eli Goz, Chen Brestel, Shir Bar, Maya Atar, and Eldad Elnekave. 2018. Textray: Mining clinical reports to gain a broad understanding of chest X-rays. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 553–561.

[22] Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. BioBERT: A pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics* 36, 4 (2020), 1234–1240.

[23] Jingshu Liu, Zachariah Zhang, and Narges Razavian. 2018. Deep EHR: Chronic disease prediction using medical notes. *arXiv preprint arXiv:1808.04928* (2018).

[24] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013).

[25] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*. 3111–3119.

[26] Maram Mahmoud A Monshi, Josiah Poon, and Vera Chung. 2020. Deep learning in generating radiology reports: A survey. *Artificial Intelligence in Medicine* (2020), 101878.

[27] Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. 1532–1543.

[28] Ewoud Pons, Loes MM Braun, MG Myriam Hunink, and Jan A Kors. 2016. Natural language processing in radiology: A systematic review. *Radiology* 279, 2 (2016), 329–343.

[29] Martin F Porter. 2001. Snowball: A language for stemming algorithms.

[30] Guergana K Savova, James J Masanz, Philip V Ogren, Jiaping Zheng, Sunghwan Sohn, Karin C Kipper-Schuler, and Christopher G Chute. 2010. Mayo clinical Text Analysis and Knowledge Extraction System (cTAKES): Architecture, component evaluation and applications. *Journal of the American Medical Informatics Association* 17, 5 (2010), 507–513.

[31] Elyne Scheurwegs, Kim Luyckx, Léon Luyten, Walter Daelemans, and Tim Van den Bulcke. 2016. Data integration of structured and unstructured sources for assigning clinical codes to patient stays. *Journal of the American Medical Informatics Association* 23, e1 (2016), e11–e19.

[32] Mike Schuster and Kuldip K Paliwal. 1997. Bidirectional recurrent neural networks. *IEEE transactions on Signal Processing* 45, 11 (1997), 2673–2681.

[33] Merlijn Sevenster, Rob Van Ommering, and Yuechen Qian. 2012. Automatically correlating clinical findings and body locations in radiology reports using MedLEE. *Journal of digital imaging* 25, 2 (2012), 240–249.

[34] Seyedmostafa Sheikhalishahi, Riccardo Miotto, Joel T Dudley, Alberto Lavelli, Fabio Rinaldi, and Venet Osmani. 2019. Natural language processing of clinical notes on chronic diseases: systematic review. *JMIR medical informatics* 7, 2 (2019), e12239.

[35] Bonggun Shin, Falgun H Chokshi, Timothy Lee, and Jinho D Choi. 2017. Classification of radiology reports using neural attention models. In *2017 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 4363–4370.

[36] Petronella Cornelia Gerarda Simons, A Algra, MF Van de Laak, DE Grobbee, and Y Van der Graaf. 1999. Second manifestations of ARTerial disease (SMART) study: Rationale and design. *European journal of epidemiology* 15, 9 (1999), 773–781.

[37] Akshay Smit, Saahil Jain, Pranav Rajpurkar, Anuj Pareek, Andrew Y Ng, and Matthew P Lungren. 2020. CheXbert: Combining automatic labelers and expert annotations for accurate radiology report labeling using BERT. *arXiv preprint arXiv:2004.09167* (2020).

[38] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15, 1 (2014), 1929–1958.

[39] Ricky K Taira, Stephen G Soderland, and Rex M Jakobovits. 2001. Automatic structuring of radiology free-text reports. *Radiographics* 21, 1 (2001), 237–245.

[40] Erik-Jan van Kesteren. 2020. Vankesteren/firatheme: Firatheme version 0.2.1. (Jan 2020). https://doi.org/10.5281/zenodo.3604681

[41] Yanshan Wang, Liwei Wang, Majid Rastegar-Mojarad, Sungrim Moon, Feichen Shen, Naveed Afzal, Sijia Liu, Yuqun Zeng, Saeed Mehrabi, Sunghwan Sohn, et al. 2018. Clinical information extraction applications: A literature review. *Journal of biomedical informatics* 77 (2018), 34–49.

[42] Zhuoran Wang, Anoop D Shah, A Rosemary Tate, Spiros Denaxas, John Shawe-Taylor, and Harry Hemingway. 2012. Extracting diagnoses and investigation results from unstructured text in electronic health records by semi-supervised machine learning. *PLoS One* 7, 1 (2012), e30412.

[43] David A Wood, Jeremy Lynch, Sina Kafiabadi, Emily Guilhem, Aisha Al Busaidi, Antanas Montvila, Thomas Varsavsky, Juveria Siddiqui, Naveen Gadapa, Matthew Townend, et al. 2020. Automated labelling using an attention model for radiology reports of MRI scans (ALARM). *arXiv preprint arXiv:2002.06588* (2020).

[44] Xiao Wu, Yuzhe Zhao, Dragomir Radev, and Ajay Malhotra. 2020. Identification of patients with carotid stenosis using natural language processing. *European Radiology* (2020), 1–9.

[45] Haotian Xu, Ming Dong, Dongxiao Zhu, Alexander Kotov, April Idalski Carcone, and Sylvie Naar-King. 2016. Text classification with topic-based word embedding and convolutional neural networks. In *Proceedings of the 7th ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*. 88–97.

[46] Keyang Xu, Mike Lam, Jingzhi Pang, Xin Gao, Charlotte Band, Pengtao Xie, and Eric Xing. 2018. Multimodal machine learning for automated ICD coding. *arXiv preprint arXiv:1810.13348* (2018).

[47] John Zech, Margaret Pain, Joseph Titano, Marcus Badgeley, Javin Schefflein, Andres Su, Anthony Costa, Joshua Bederson, Joseph Lehar, and Eric Karl Oermann. 2018. Natural language–based machine learning models for the annotation of clinical radiology reports. *Radiology* 287, 2 (2018), 570–580.