

HITSZ-ICRC: A Report for SMM4H Shared Task 2020-Automatic Classification of Medications and Adverse Effect in Tweets

Xiaoyu Zhao, Ying Xiong, Buzhou Tang

Department of Computer Science, Harbin Institute of Technology (Shenzhen), China
zhaoxiaoyux@126.com, {xiongying0929, tangbuzhou}@gmail.com

Abstract

This is the system description of the Harbin Institute of Technology Shenzhen (HITSZ) team for the first and second subtasks of the fifth Social Media Mining for Health Applications (SMM4H) shared task in 2020. The first task is automatic classification of tweets that mention medications and the second task is automatic classification of tweets in English that report adverse effects. The system we propose for these tasks is based on bidirectional encoder representations from transformers (BERT) incorporating with knowledge graph and retrieving evidence from online information. Our system achieves an F1 of 0.7553 in task 1 and an F1 of 0.5455 in task 2.

1 Introduction

With greatly increasing number of social media users, plenty of personal information has been posted in social networking websites such as Twitter, including massive data about health and medicine. Social media has attracted more and more attention from researchers and industrials in the health and medicine domain. To utilize the imbalanced and informal expressions of clinical concepts, a system should be able to alleviate the noise of social media posts and detect medical mentions which requires professional knowledge. In 2020, the health language processing lab at University of Pennsylvania organized the Social Media Mining for Health Applications (SMM4H) shared task about using social media mining for health monitoring, which includes 5 subtasks. We participate in the first two subtasks: (1) Automatic classification of tweets that mention medications, and (2) Automatic classification of tweets in English that report adverse effects (AEs). In this report, we briefly introduce our system developed for the two subtasks. The system is based on pre-trained language model BERT (Devlin, Chang, Lee, & Toutanova, 2018) with external knowledge.

2 Data and Pre-processing

The 2020 SMM4H organizers provide a dataset of tweets with a label indicating whether the tweets mention medications for task1 (1-Yes, 0-No), and a dataset of tweets with a label indicating whether the tweets mention AEs for task2. Each tweet includes text, a date and an identified ID. Table 1 describes the statistics of the two individual datasets, where #1 and #0 denote the numbers of tweets that mention medications and not respectively, #all denotes the total number of tweets, and NA denotes that we do not know the number of those tweets.

Task	Dataset	#1	#0	#all
Task1	training set	181	69,091	69,272
	test set	NA	NA	29,687
Task2	training set	2,374	23,298	25,672
	test set	NA	NA	4,659

Table 1: Distribution of label over the training and test datasets of task 1 and task 2.

This work is licensed under a Creative Commons Attribution 4.0 International Licence. Licence details:
<http://creativecommons.org/licenses/by/4.0/>.

For both tasks, we pre-process all tweets via removing the usernames, URLs, emoticons, non-ASCII characters and repeated punctuations.

3 System Description

Our system for both task 1 and task 2 is based on BERT, which is used to represent tweet. Below we describe in detail the methods for the two tasks one by one.

3.1 Methods for task1

In this task, we first deploy BERT directly (Shuai Chen, 2019), and then extend it to BERT_Med by introducing potential medications from DrugBank¹. The architecture of BERT_Med is shown in Fig. 1, where FC is a fully-connected layer. Given a piece of tweet, we first calculate its similarities with medications, and then select those medications with similarity less than 0.1 as its next tweet, finally input the tweet and the selected medications into BERT. The similarity between tweets and medications is measured by cosine using TF-IDF(Term Frequency-Inverse Document Frequency) as their representations.

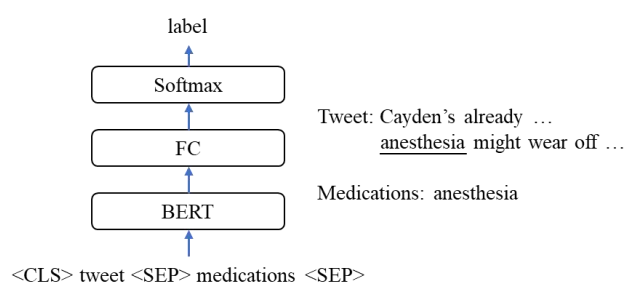


Figure 1: BERT with potential drugs (BERT_Med)

3.2 Methods for task2

For task 2, we propose the following two models based on BERT as shown in Fig. 2.

Strategy 1: To extract the internal information of AEs in tweets, we add the potential AEs corresponding to the drug mentions detected in the tweets. Drug mentions in tweets are detected by exact matching on Drugbank and the drug dictionary (Azadeh N, Abeer S, O'Connor Karen, et al. 2015), and the AEs come from MedlinePlus². Similar to BERT_MED in Fig. 1, we treat the potential AEs as the next tweet of the given tweets and input them into BERT. In addition, we add a self-attention module (Ashish Vaswani, Noam Shazeer, Niki Parmar and et al. 2017) after BERT to redistribute representations of the tweet and potential AEs.

Strategy 2: To further integrate external knowledge into strategy 1, we introduce medication representation learnt from MeSH³ by TransE (Bordes A, Usunier N, Garcia-Duran A, et al. 2013) as shown in Fig. 2. A similar self-attention module is employed to extract information from the concatenated entity embeddings learnt by TransE. Finally, we feed combination of the final representation of the texts and the medication entities into a linear transformation layer for the final prediction.

¹ <https://www.drugbank.ca/>

² <https://medlineplus.gov/>

³ <https://www.nlm.nih.gov/mesh/meshhome.html>

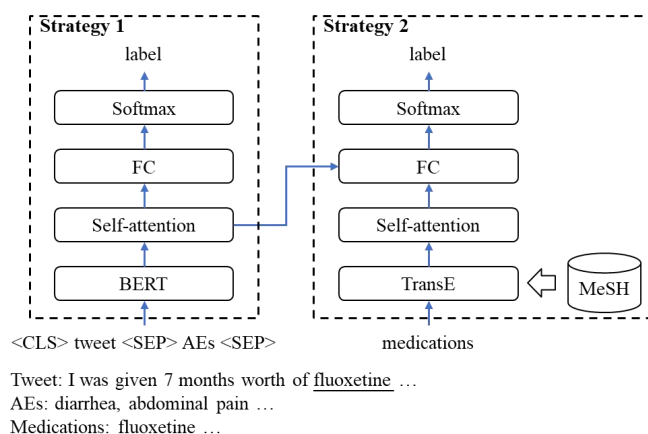


Figure 2: Model architectures in task 2.

4 Experiments

For task1, we set batch size to 64, learning rate to $2e-5$ and sequence length to 64 when training models. For task2, the batch size and sequence length of BERT are the same as the ones of task1 while the learning rates of BERT and attention layer are set to $1e-5$ and $5e-5$ respectively. We train all models using 5-fold cross validation on the training set.

5 Results

Task	System	Training set			Test set		
		F1	P	R	F1	P	R
Task 1	BERT	0.8436	0.9310	0.7714	0.7553	0.8182	0.7013
	BERT_Med	0.6667	0.5246	0.9143	NA	NA	NA
Task 2	BERT_Strategy1	0.6877	0.6501	0.7299	0.5190	0.4393	0.6340
	BERT_Strategy2	0.6802	0.6823	0.6741	0.5254	0.4460	0.6392
	BERT_Merged	0.6701	0.6177	0.7321	0.5455	0.4701	0.6495

Table 2: Results on training and test sets for task 1 and task 2.

The results of our system provided by the organizers are listed in Table 2, where NA denotes that we do not submit results as internet problem, BERT_Merged is the model that merges results of BERT_Strategy1 and BERT_Strategy2 by voting. Our system achieves the highest F-score of 0.7553 on task1 and 0.5455 on task2. For task 1, on the training set, BERT_Med achieves much higher recall, but much lower precision than BERT. For task2, BERT_Strategy1 and BERT_Strategy2 perform roughly the same on the training and test sets. BERT_Merged performs slightly worse than them on the training set, but a little higher on the test set. The difference between performances of our system on the training set and test set is very large.

Reference

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. *Bert: Pre-training of deep bidirectional transformers for language understanding*. arXiv preprint arXiv:1810.04805.
- Shuai Chen, Yuanhang Huang, Xiaowei Huang, Haoming Qin, Jun Yan, Buzhou Tang. 2019. *HITSZ-ICRC: A report for SMM4H shared task 2019-automatic classification and extraction of adverse effect mentions in tweets*. Proceedings of the Fourth Social Media Mining for Health Applications (# SMM4H) Workshop & Shared Task. pp. 47-51.
- Azadeh N, Abeed S, O'Connor Karen, et al. 2015. *Pharmacovigilance from social media: mining adverse drug reaction mentions using sequence labeling with word embedding cluster features*. Journal of the American Medical Informatics Association. (3):3.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser and Illia Polosukhin. 2017. Attention Is All You Need. arXiv preprint arXiv: 1706.03762.
Bordes A, Usunier N, Garcia-Duran A, et al. 2013. *Translating embeddings for modeling multi-relational data*. Proc of NIPS. Cambridge, MA: MIT Press, 2013:2787-2795