Detecting Potential Adverse Drug Reactions from Health-Related Social Networks

Bo Xu, Hongfei Lin(✉), Mingzhen Zhao, Zhihao Yang, Jian Wang, and Shaowu Zhang

School of Computer Science and Technology, Dalian University of Technology, Dalian 116024, Liaoning, China
hflin@dlut.edu.cn

Abstract. In recent years, adverse drug reactions have drawn more and more attention from the public, which may lead to great damage to the public health and cause massive economic losses to our society. As a result, it becomes a great challenge to detect the potential adverse drug reactions before and after putting drugs into the market. With the development of the Internet, health-related social networks have accumulated large amounts of users' comments on drugs, which may contribute to detect the adverse drug reactions. To this end, we propose a novel framework to detect potential adverse drug reactions based on health-related social networks. In our framework, we first extract mentions of diseases and adverse drug reactions from users' comments using conditional random fields with different levels of features, and then filter the indications of drugs and known adverse drug reactions by external biomedical resources to obtain the potential adverse drug reactions. On the basis, we propose a modified Skip-gram model to discover associated proteins of potential adverse drug reactions, which will facilitate the biomedical experts to determine the authenticity of the potential adverse reactions. Extensive experiments based on DailyStrength show that our framework is effective for detecting potential adverse drug reactions from users' comments.

Keywords: Adverse drug reactions · Health-related social network · ADRs

1 Introduction

Adverse drug reactions (ADRs) have drawn more and more attention from the public, which may not only lead to serious physical injuries, but also cause great economic losses. It is estimated that each year about 2 million patients in the United States experience serious ADRs by using marketed drugs, resulting in more than 100,000 deaths. ADRs are considered to be the fourth leading cause of death [1], and about $136 billion is spent on treating ADRs in the United State every year [2]. Therefore, research on detecting adverse drug reactions has been studies for years. For example, Rahmani et al. [3] proposed a novel network-based approach to predict the ADRs by modeling the interactions among different drugs. Casillas et al. [4] present a hybrid system utilizing a self-developed morpho-syntactic and semantic analyzer for medical texts in Spanish. These studies detect ADRs mainly based on extracting drug-drug

C.-Y. Lin et al. (Eds.): NLPCC-ICCPOL 2016, LNAI 10102, pp. 523–530, 2016.
DOI: 10.1007/978-3-319-50496-4_45

interactions or drug reaction events from existing literatures. However, it remains a great challenge to automatically detect potential ADRs, which have not been verified.

In recent years, health-related social networks have attracted much attention from the public, which accumulate large amounts of users' comments about drugs. These comments contain a great deal of information related to potential ADRs. To capture the information, Leaman et al. [2] use dictionary-based method to recognize the mentions of ADRs, and achieve promising results, which is one of the earliest researches to mine the relationships between drugs and adverse reactions from health-related social networks. Dai et al. [5] focus on the recognition of ADRs from twitter in terms of feature engineering by extracting various features and examining the performance of different combinations of features. Since tweets are short texts and contain many colloquial expressions, there exist lots of noises in the corpus, producing negative impact on their results. To deal with the problem, Yates and Goharian [6] annotate 2500 pieces of users' comments on five drugs on breast cancer to find the patterns of the mentions of ADRs in social networks. These studies show that comments on social networks are potentially useful for detecting ADRs.

In this paper, we propose a novel framework to detect potential ADRs from health-related social networks. In the framework, we adopt the conditional random fields with different combinations of features to recognize the mentions of diseases and ADRs. After that, we filter the indications of drugs and known ADRs to obtain the potential ADRs. On the basis, we also seek to find the associated proteins for the potential ADRs, which can link the drugs to the potential adverse reactions, and give adequate evidence for ADR verifications.

2 Detecting Potential ADRs on Health-Related Social Networks

Our framework is designed to detect potential adverse drug reactions from health-related social networks, and discover the associated proteins on ADRs to provide the most likely evidence chains for practitioners to verify the ADRs. Overall, there are three modules in our framework, including the data acquisition module, potential ADRs detecting module and the associated protein recognition modules. We will introduce each module in details in the following subsections.

2.1 Data Acquisition Module

To obtain the data from the DailyStrength, we use Scrapy (http://scrapy.org/) program to crawl the comments of drug takers. Scrapy, implemented in Python, is an open source and collaborative framework for acquiring the data from the Internet. We obtain the comments mainly on the boards of drugs with most posts for further processing. After obtaining the drug posts, we preprocess the data by removing special characters and noisy posts with only a few words. It should be noted that since our method is general, it can also be applied for other health-related social networks, such as MedHelp or Ask a Patient.

2.2 Potential ADRs Detecting Module

There are three steps in detecting potential ADRs based on the users' comments, including named entity recognition for disease and ADRs, filtering the indications and filtering the known ADRs. We introduce the steps in details as follows.

As the first step, we recognize named entities for diseases and ADRs. Users' comments from health-related social network contain both ADR names and disease names, we recognize both names in the phase of named entity recognition.

We take the problem of extracting the names from users' comments as a sequence labeling problem, solved using conditional random field (CRF) model. To apply the CRF model for extracting names of diseases and ADRs, we first define some features to discriminate the mentions of drugs and diseases from other elements. Overall, we define two kinds of features, word-level features and dictionary-based features.

For word-level features, we define three features, the original word feature (OW), the stemmed word feature (SW) and part-of-speech word feature (POS). We use Stanford POS Tagger [7] to generate the POS for words. For dictionary-based features, we extract domain-specific features to improve the performance of CRF by selecting the indications and ADRs to generate a disease and ADR name dictionary from SIDER [8], which is a standard drug adverse reactions database. We define four features in this set, namely SIDER-based feature (SF), first-word feature (FF), last-word feature (LF) and single-word feature (SinF). SIDER-based feature can be extracted based on whether a word exists in the disease and ADR name dictionary SIDER. First-word feature measures whether a word is at the first position of some disease names or ADR names in the dictionary. Last-word feature measures whether a word is at the last position of names in the dictionary. Single-word feature is determined by considering whether a word is a disease name or ADR name. We train CRF model with all the defined features to extract the mentions of disease names and ADR names.

As the second step, we filter drug indications. In users' comments, there exists a mixture of adverse reactions and drug indications. Indications refer to the corresponding symptoms of disease a drug treats, which would be of less use for detecting the potential ADRs. Therefore, we seek to filter indications of drugs. Since the users' comments are organized by drug names from DailyStrength, we can easily learn which drug some users' comments belong to, and use external resources to filter the indications.

To find the indications of drugs, we resort to the DrugBank [9] database and the Semantic MEDLINE database (SemMedDB) [10] to filter the indications. We first adopt DrugBank to obtain indications with respect to certain drugs, and then use SemMedDB to filter other indications. SemMedDB contains a large amounts of triples in the style of (*subject*, *predicate*, *object*) extracted from MEDLINE by SemRep [11], a semantic interpreter of biomedical text. If a triple <*drug_i*, *TREATS*, *symptom_j*> exists in the SemMedDB, we know *symptom_j* is one indication of *drug_i*.

As the third step, we filter known ADRs. After filtering the indications, we obtain the mentions of ADRs, some of which have been recorded in official instructions. We take the recorded ADRs as known ADRs. We filter the known ADRs using SIDER database, which is an official released adverse drug reaction database.

2.3 Associated Protein Recognition Modules

In the field of clinical medicine, verifying ADRs require much time and efforts by large amounts of clinical trials and observations. Therefore, to facilitate the verifications, we adopt text mining techniques to discover the associated proteins, which refer to the effected proteins by taking a certain drug, and may cause some adverse reactions. Generally, associated protein can be taken as an evidence of adverse drug reactions, and provide much help for medical experts to verify the potential ADRs.

To recognize associated proteins, we propose a modified Skip-gram model to generate distributed representations of drugs, proteins and potential ADRs, and measuring the similarity degrees among them. The Skip-gram model is proposed by Google [12, 13], which generates distributed representations for words in the training corpus, capturing the syntactic and semantic relationships among them [13].

We use citations in MEDLINE with annotated Medical Subject Headings (MeSH) terms to train the distributed representations. The MeSH terms can reflect the topics of each article, and the co-occurrences of the MeSH terms in one article indicate the correlations among the entities. Therefore, we train the Skip-gram model based on the MeSH term sets of the articles about certain drugs. Since the original Skip-gram model is trained on corpus using slide windows and the MeSH term sets are out-of-order, we modify the Skip-gram model by transforming the sliding window to document window, taking every MeSH term as a single word.

After generating distributed vectors for every MeSH term, we define the association function for each triple of the drug, one associated protein and one potential ADR (*d, p, a*) as follow.

$$f(d,p,a) = \frac{sim(d,p) + sim(p,a)}{1 + |sim(d,p) - sim(p,a)|} \tag{1}$$

where d is the drug, a is one potential adverse reaction, p is associated protein. $sim(x, y)$ measures the similarity between entity x and entity y. intuitively, if $sim(d, p) + sim(p, a)$ is larger, the protein p is more likely to be the associated protein between drug d and the potential ADR a. To avoid that $sim(d, p)$ or $sim(p, a)$ is too large to impact the result, we normalize the function using the smoothing factor as the denominator of Eq. (1). The function f can be used to measure the association degree among the items in the triple (d, p, a). For all the potential ADRs of the drug d, we sort the associated proteins based on the association function $f(d, p, a)$ and take the top-k proteins as the final associated proteins.

3 Experiments and Result Analysis

3.1 Datasets

In our experiments, we crawl users' comments before June 2, 2014 from health-related social network DailyStrength with respect to 50 most focused drugs. There are totally 600,237 pieces of comments in 1075 health-related topics. We also use the annotated MeSH sets of article citations on certain drugs, containing more than 22 million

MEDLINE citations before the year of 2013 and extract the MeSH sets from them as the train data for training the modified Skip-gram model. We set the dimension of the vectors to be 100 for relatively good performance. We conduct three groups of experiments to evaluate the performance of our framework. The first experiment is conducted to examine the performance of CRF based named entity recognition with different subsets of features. The second experiment is designed to verify the detected ADRs based on the literatures. The third experiment is to find the associated proteins with respect the potential ADRs.

3.2 Performance on Recognizing Mentions of Diseases and ADRs

To train the CRF model for recognizing names of diseases and ADRs, we annotated 2000 pieces of users' comments in our experiments. For every combination of features, we evaluate the performance of the CRF model on accuracy, recall and F1 scores using 10-fold cross evaluation. Table 1 shows the results for CRF with different feature combinations. Compared with the original word features, the set with all of the defined features achieves the best performance, which achieves the best F1 score 83%, while the F1 score is 82% when using only the word-level feature set. The performance is improved by adding the dictionary-based feature set.

Table 1. Performance on recognizing mentions by the CRF model

Feature combinations	Recall	Accuracy	F1
SW + POS	0.76	0.87	0.81
OW + POS	0.74	0.88	0.80
OW + POS + SW	0.76	0.88	0.82
OW + POS + SW + SF	0.77	0.88	0.82
OW + POS + SW + SF + FF	0.77	0.87	0.82
OW + POS + SW + SF + FF + LF	0.77	0.87	0.82
ALL	0.78	0.87	0.83

3.3 Performance on Potential ADRs Detection

In this section, we take three drugs as examples to compare the detected top-10 ADRs by our method with those by Leaman's Method [2] in Table 2. In the table, "-" stands for the indication of drugs, "+" stands for the known ADRs and "*" stands for the potential ADRs. The Sim. scores measure the confidence degree to verify the potential ADRs with respect to the corresponding drug by each method. From the table, we can find that the detected ADRs by these two methods are highly correlated, which indicates our method is as effective as the state-of-the-art work.

On this basis, we filter the indications and ADRs using some domain-specific resources. We totally extract 993 item mentions, 231 of which can be verified as indications using semantic MEDLINE and 34 of which can be verified as indications using DrugBank. For the recognized potential ADRs, 240 of which can be verified using the SIDER database and 488 of which can be taken as potential ADRs. From the

Table 2. Comparison with Leaman's results

Drug names	Leaman's results [2]		Our results	
	Mentions	Sim.	Mentions	Sim.
Carbamazepine	somnolence or fatigue	12.3%	seizures-	29.6%
	allergy	5.2%	not as effective*	23.8%
	weight gain	4.1%	pain-	17.0%
	rash	3.5%	rash-	10.2%
	depression	3.2%	sleepy*	4.4%
	dizziness	2.4%	effect increased*	3.9%
	tremor/spasm	1.7%	weight gain abnormal*	3.2%
	headache	1.7%	dizziness+	2.7%
	appetite increased	1.5%	headache-	2.7%
	nausea	1.5%	nausea+	2.7%
Trazodone	somnolence or fatigue	48.2%	insomnia-	18.1%
	nightmares	4.6%	anxiety*	12.8%
	insomnia	2.7%	not as effective*	10.9%
	addiction	1.7%	wakefulness*	10.4%
	headache	1.6%	sleepy*	7.6%
	depression	1.3%	nightmare-	7.3%
	hangover	1.2%	hangover effect*	4.8%
	anxiety attack	1.2%	feeling high*	4.4%
	panic reaction	1.1%	drowsiness-	3.9%
	dizziness	0.9%	headache+	3.4%
Ziprasidone	somnolence or fatigue	20.3%	not as effective*	33.7%
	dyskinesia	6.0%	sleepy*	15.2%
	mania	3.7%	anxiety+	9.6%
	anxiety attack	3.5%	mania-	6.3%
	weight gain	3.2%	weight gain abnormal*	4.5%
	depression	2.4%	hallucination*	4.3%
	allergic reaction	1.9%	suicide*	3.9%
	dizziness	1.2%	feeling high*	2.9%
	panic reaction	1.2%	effect increased*	2.5%

table, we also find that drug takers are likely to comment on the ADRs not presented in the instructions. That is to say, if the drug instruction already tells that the drug may cause some kinds of ADRs, the takers would consider the presented ADRs as normal reactions before taking the drug. Otherwise, the drug takers tend to resort to the health-related social network for help.

After the extraction, we filter the indications and known ADRs. In the known ADRs, 71 of them can be verified by Semantic MEDLINE, which have not been recorded in SIDER.

3.4 Associated Proteins for Potential ADRs

We consider ADRs not existed in the SIDER and without evidences from semantic MEDLINE as potential ADRs, and attempt to find associated proteins using the

distributed entity representations. We list the associated proteins between *trazodone* and *anxiety* in Table 3, and try to verify their association from existing literatures. From Gingrich's work [14], we can find that serotonin receptors are related to anxiety. Goldman et al. [15] indicate serotonin transporter is related to anxiety. As a special serotonin transporter, we can infer that serotonin plasma membrane transport proteins are related to anxiety. Shishkina et al. [16] indicate that adrenergic receptors are related to anxiety. On the other hand, trazodone is an antidepressant of the serotonin antagonist and reuptake inhibitor class and it has the alpha-adrenergic blocking property. So in the five detected associated proteins, we can obtain three triples, <*trazodone*, *serotonin receptors*, *anxiety*>, <*trazodone*, *serotonin plasma membrane transport proteins*, *anxiety*> and <*trazodone*, *adrenergic receptors*, *anxiety*>, which will help the biomedical experts determine the relationships between trazodone and anxiety.

Table 3. The associated proteins of trazodone and anxiety

Associated proteins	$f(d,p,a)$
receptors, serotonin	0.83
5-hydroxytryptophan	0.76
serotonin plasma membrane transport proteins	0.75
receptors, adrenergic	0.75
receptor, serotonin, 5-ht1a	0.74

4 Conclusion and Future Work

In this paper, we propose a novel framework to detect potential adverse drug reactions from health-related social networks. In the framework, we first extract mentions of diseases and ADRs using CRF with different features. Then, we filter indications of drugs and known ADRs with the help of biomedical databases, including SIDER, DrugBank and Semantic MEDLINE, to obtain the potential ADRs. Finally, to facilitate verifications of potential ADRs, we propose a modified Skip-gram model to discover the associated proteins between a certain drug and its corresponding potential ADRs by generating distributed representations of biomedical entities. Experimental results show the effectiveness of our framework to detect the potential ADRs from social network DailyStrength. In our future work, we will attempt to generate representations for other biomedical entities, not limited to MeSH, and develop other effective method to find associations for potential ADRs beyond associated proteins.

Acknowledgements. This work is partially supported by grant from the Natural Science Foundation of China (Nos. 61277370, 61402075, 61572102, 61632011, 61602078, 61572098), Natural Science Foundation of Liaoning Province, China (Nos. 201202031, 2014020003), State Education Ministry and The Research Fund for the Doctoral Program of Higher Education (No. 20090041110002), the Fundamental Research Funds for the Central Universities. The 12th five year national science and technology supporting programs of China under Grant No. 2015BAF20B02.

References

1. Giacomini, K.M., Krauss, R.M., Roden, D.M., Eichelbaum, M., Hayden, M.R., Nakamura, Y.: When good drugs go bad. Nature **446**(7139), 975–977 (2007)
2. Leaman, R., Wojtulewicz, L., Sullivan, R., Skariah, A., Yang, J., Gonzalez, G.: Towards internet-age pharmacovigilance: extracting adverse drug reactions from user posts to health-related social networks. In: Paper Presented at the Proceedings of the 2010 Workshop on Biomedical Natural Language Processing (2010)
3. Rahmani, H., Weiss, G., Méndez-Lucio, O., Bender, A.: ARWAR: a network approach for predicting adverse drug reactions. Comput. Biol. Med. **68**, 101–108 (2016)
4. Casillas, A., Pérez, A., Oronoz, M., Gojenola, K., Santiso, S.: Learning to extract adverse drug reaction events from electronic health records in Spanish. Expert Syst. Appl. **61**, 235–245 (2016)
5. Dai, H.J., Touray, M., Jonnagaddala, J., Syed-Abdul, S.: Feature engineering for recognizing adverse drug reactions from Twitter posts. Information **7**(2), 27 (2016)
6. Yates, A., Goharian, N.: ADRTrace: detecting expected and unexpected adverse drug reactions from user reviews on social media sites. In: Serdyukov, P., Braslavski, P., Kuznetsov, S.O., Kamps, J., Rüger, S., Agichtein, E., Segalovich, I., Yilmaz, E. (eds.) ECIR 2013. LNCS, vol. 7814, pp. 816–819. Springer, Heidelberg (2013). doi:10.1007/978-3-642-36973-5_92
7. Toutanova, K., Klein, D., Manning, C.D., Singer, Y.: Feature-rich part-of-speech tagging with a cyclic dependency network. In: Paper presented at the Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology-Volume 1 (2003)
8. Kuhn, M., Campillos, M., Letunic, I., Jensen, L.J., Bork, P.: A side effect resource to capture phenotypic effects of drugs. Mol. Syst. Biol. **6**(1), 343 (2010)
9. Wishart, D.S., Knox, C., Guo, A.C., Shrivastava, S., Hassanali, M., Stothard, P., Woolsey, J.: DrugBank: a comprehensive resource for in silico drug discovery and exploration. Nucleic Acids Res. **34**(suppl 1), D668–D672 (2006)
10. Wishart, D.S., Knox, C., Guo, A.C., Cheng, D., Shrivastava, S., Tzur, D., Hassanali, M.: DrugBank: a knowledgebase for drugs, drug reactions and drug targets. Nucleic Acids Res. **36**(suppl 1), D901–D906 (2008)
11. Rindflesch, T.C., Fiszman, M.: The interaction of domain knowledge and linguistic structure in natural language processing: interpreting hypernymic propositions in biomedical text. J. Biomed. Inf. **36**(6), 462–477 (2003)
12. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. arXiv preprint arXiv:1301.3781 (2013)
13. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: Paper presented at the Advances in Neural Information Processing Systems (2013)
14. Gingrich, J.A.: Mutational analysis of the serotonergic system: recent findings using knockout mice. Curr. Drug Targets-CNS Neurol. Dis. **1**(5), 449–465 (2002)
15. Goldman, D., Oroszi, G., Omalley, S., et al.: COMBINE genetics study: the pharmacogenetics of alcoholism treatment response: genes and mechanisms. J. Stud. Alcohol Suppl. **66**(15), 56–64 (2005). discussion 33
16. Shishkina, G., Kalinina, T., Dygalo, N.: Attenuation of α 2A-adrenergic receptor expression in neonatal rat brain by RNA interference or antisense oligonucleotide reduced anxiety in adulthood. Neuroscience **129**(3), 521–528 (2004)