Proceedings of the Thirtieth International Florida Artificial Intelligence Research Society Conference

Unsupervised Aspect Term Extraction in Online Drugs Reviews

Diana C. Cavalcanti and Ricardo B. C. Prudêncio
Centro de Informática, Universidade Federal de Pernambuco, Cidade Universitária - 50740-560, Recife/PE, Brazil
{dcc2,rbcp}@cin.ufpe.br

Abstract

Aspect mining in drugs reviews has focused on extracting relevant information such as adverse reactions, efficacy of a drug, symptoms and conditions of patients. In our work, a new unsupervised and knowledge-based method is proposed for extracting aspects in drug reviews. The proposed solution is based on linguistic features, more specifically dependency paths in the syntactic tree of a review. The quality of the dependency path rules was investigated in a number of experiments in review corpora associated to three different diseases. Promising results were achieved compared to previous work.

Introduction

Sentiment and opinion about symptoms, treatments or medicines expressed in online media provide new opportunities for researches on opinion mining and sentiment analysis (Denecke and Deng 2015). Sentiment in medical context can be seen as a reflection of the health status of a patient, which can be good, bad or normal in some time interval. Expressed terms, like *severe pain*, can be a good indication of the health and quality of life of a patient. Descriptions of the use of a medicine may be related to symptoms presented before ingestion, such as *anxiety co-morbid*, *high blood pressure* or about adverse reactions, such as *extreme weight loss*.

Adverse Drug Reactions (ADR) have been commonly reported by patients in drug reviews. ADRs are harmful reactions caused by medication intake resulting in an intervention related to the use of the product, which provides risk in future use or specific treatment, change in dosage or even the withdrawal of product market (Yang and Yang 2015). The activities relating to the detection, assessment, understanding and prevention of adverse effects related to drugs is known as pharmacovigilance or drug safety monitoring. Pharmacovigilance starts during clinical trials of a drug and continued after released it for consumption (Gosal 2015).

Due to various limitations in clinical trials, it is not possible to fully evaluate the consequences of using a particular drug before being released (Gosal 2015; Cheng et al. 2014). An ADR may not be detected before the product go to the market and can take some time after its sale to track new ADRs and relate them to the drug's label (Sampathkumar, Chen, and Luo 2014).

Opinion mining is described as the process of detecting opinions and opinionated aspects in subjective texts from large volumes of structured or unstructured texts using computational methods (Veloso and Jr. 2007; Pang and Lee 2008; Harpaz et al. 2014). Many scientists have focused their research on the development of mining techniques in medical and pharmaceutical texts from publicly available data on the web. Specifically in pharmacovigilance, drug manufacturers can benefit from opinion mining since particular adverse reactions to a drug can be traced more quickly from public repositories or posts in social networks.

In this work, we focused on aspect-based opinion mining (Liu 2012) in drug reviews. The aim is to identify in a drug review fragments of text that can be associated to specific aspects of interest like, adverse reactions, effectiveness, patient conditions, among others. In general, we can distinguish in the literature two approaches for this task: (1) machine learning, in which sequential labeling algorithms (like Hidden Markov Models) are used to label the words in the review; and (2) knowledge-based systems, in which relevant parts of the review are extracted by the use of linguistic rules. There are advantages and limitations of each approach. We focused on the knowledge-based approach with linguistic techniques, since there is no need for a labeled training corpus, which can be very demanding in practice.

A new method for extracting aspects in drug reviews is proposed based on dependency paths identified in the syntactic tree of a review. Dependency path rules have been successfully adopted in other domains (Bancken, Alfarone, and Davis 2014). In our proposal, we derived new rules specific to the drug review domain and investigated their performance through experiments in three datasets related to different diseases, labeled with relevant aspects. The results revealed a gain in performance (in terms of F-Measure) compared to previous work.

The current work filled in a gap in the literature by investigating and proposing effective knowledge-based and unsupervised methods for aspect extraction in drug reviews. We can mention the following contributions:

- Proposal of a new method for extracting aspects in drug reviews based on syntactic dependency paths, which shown to have a good predictive performance;
- Investigation of previous knowledge-based systems that

use linguistic features for aspect extraction and their adequacy to the drug review domain;

- Production of a corpus of drug reviews labeled by aspects, which can be adopted for new experiments.

The remainder of this paper is organized as follows. Section II provides a brief introduction on opinion mining in drug reviews. In Section III we present the proposed approach. Section IV presents the results of our experiments. Finally, Section V concludes the paper with a discussion of our results, along with recommendations for future work.

Opinion mining in drugs reviews

Recent studies have focused on mining reviews of drugs (available in social networks, forums, web,...) in order to provide useful information for pharmacovigilance. Two common tasks can be identified in the literature of opinion mining in drug reviews: (1) classification of reviews; and (2) extraction of opinion aspects.

Concerning the first task, a review can be automatically classified according to its sentiment about the drug (usually positive, negative or neutral) (Na et al. 2012; Na and Kyaing 2015). Sentiment classification can be useful for filtering relevant reviews to be inspected (e.g., reviews classified as negative have a high chance of mentioning a negative side effect). A related task is to directly classify whether a review mentions an ADR or not (Egger, Uzdilli, and Cielebak 2015; Jonnagaddala, Jue, and Dai 2016). This is more specific and focused than simply classifying the sentiment of a review.

Previous work on review classification adopted in general two categories of techniques: (1) machine learning and (2) knowledge-based approaches. In the machine learning approach, a classification model is learned from a labeled set of reviews. These works can be distinguished by the learning algorithm and the features used for classification. For instance, in (Jonnagaddala, Jue, and Dai 2016; Egger, Uzdilli, and Cielebak 2015), Support Vector Machines (SVMs) were adopted as classifier and different features like n-grams, part-of-speech (POS) tags and lexicon words were considered for classifying tweets that mention an ADR. In (Sharif et al. 2014), feature ensemble was proposed to reduce the sparsity of the feature representation of reviews. Sarker and Gonzalez (2015) adopted different classifiers (Naive Bayes, Maximum Entropy (MNB) and SVMs) and lexicon features. In (Patki and Gonzalez 2014), MNB and SVMs were adopted and the *Wordnet* lexicon has been used to expand synonyms of verbs, adjectives and nouns. A sentiment lexicon is also employed in this work.

A known disadvantage of the machine learning approach is the need of a labeled corpus for training, which can be prohibitive in some applications. Alternatively, previous work has adopted knowledge-based approaches for review classification. For instance, in (Na et al. 2012; Na and Kyaing 2015), the authors proposed a set of rules based on linguistic features to identify patterns in the reviews and then to perform the sentiment classification. The adopted rules are based on semantic dependency analysis (extraction of grammatical dependencies in the texts) and a subjective lexicon. In (Noferesti and Shamsfard 2015), the authors proposed rules for classifying the sentiment polarity of opinions previously extracted from the reviews. The classification is based on POS tagger information and a domain knowledge base.

The second task of opinion mining in drug reviews in to extract specific aspects. Hence, the aim is not only to classify whether a review mentions an ADRs, for instance, but also to extract the reaction itself or any other aspect considered as relevant (Sampathkumar, Chen, and Luo 2014). In literature, this task is called *aspect-based opinion mining*, which aims to extract the main aspects mentioned about an item or entity and to provide the classification of the opinion given on every aspect (Liu 2012). According to (Na et al. 2012; Na and Kyaing 2015; Nikfarjam et al. 2015), six aspects (see examples in Table 1) are common in drug reviews:

- **Overall:** The general opinion of a medicinal product or when the clause is not mentioned any of the other five categories of aspects;
- **Effectiveness:** Changes noted after the use of the medicine, linked to the patient's condition or disease;
- **Side effects:** The reactions that are not related to the medicine. ADR is a negative side effect;
- **Dosage:** Reports the amount, frequency or the treatment period in which the medicine was used.
- **Condition:** Corresponds to a description of the patient's condition, e.g., a disease or health problems in general.
- **Cost:** Corresponds the cost/price of a medicine.

Table 1: Opinion sentence by aspects in drugs reviews

Aspect	Opinion sentence
Overall	Adderall is overall a **good** ADD **medicine**.
Effectiveness	It **helped** me stay **focused** on any tasks.
Side effects	I have only one side effect which is **dry mouth**.
Condition	I **had** clinical **depression/anxiety** for years.
Cost	I hate that the **price** is so **high**.
Dosage	I take **30mg twice a day**

As in the classification task, previous work on aspect-based opinion mining can be split in the two categories: (1) machine learning; and (2) knowledge-based approaches. In the machine learning approach, sequential learning algorithms like Conditional Random Fields (CRFs) and Hidden Markov Models (HMMs) have been adopted specially to extract ADRs, but also other aspects (Sampathkumar, Chen, and Luo 2014). The idea is to treat aspect extraction as a sequence labeling task: each word in a review is labeled with a tag associated to an aspect and sequences of words with the same tag are extracted. Sequential learning also requires labeled corpora for training the models. The need for a large training set can be even more critical for sequential learning, since whole sequences of inputs have to be classified.

Another task is the creation of medical opinion lexicon, usually created using information of a general lexicon. In (Asghar et al. 2013), a subjectivity lexicon is proposed based

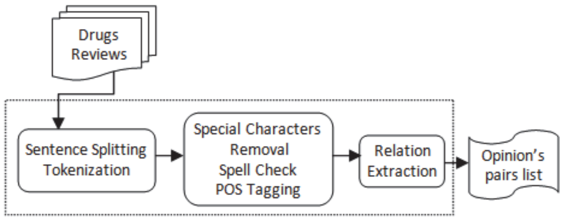

Figure 1: An overview of the proposed approach

on the corpus of drug reviews. The lexicon takes the initial medical seed list as input, expands it with SentiWordNet synonyms and antonyms, attaching polarity score with each word. In (Asghar et al. 2016), a health-related sentiment lexicon is proposed by a hybrid approach, which combines boot-strapping concepts and corpus-based strategies.

Few attempts have been identified in the literature to build knowledge-based techniques for aspect extraction in drug reviews, although such techniques are very common in the general application of aspect-based opinion mining (e.g., for product reviews). Additionally, the existing works are not completely adequate for the task. For instance, in (Noferesti and Shamsfard 2015), a knowledge-based approach is proposed for extracting text fragments in a review that express an opinion about an entity. Although this approach can be used for filtering relevant opinions in the reviews, it does not directly extract specific aspects (like ADRs). In (Na et al. 2012; Na and Kyaing 2015), text fragments are extracted from reviews with the focus on classifying sentiments. Although some of the extracted information can coincide with aspects of interest, this approach is also not focused on aspect-based extraction. So there is a gap in the literature that has to be addressed by new research.

Proposed Approach

Knowledge-based approaches which adopted linguistic rules are an interesting alternative for aspect mining. Previous techniques successfully adopted in other domains (like product reviews) can be investigated in the domain of drug reviews, thus filling in a gap identified in the literature. Obviously, adaptation of previous methods has to be addressed to fit the specific characteristics of the drug review domain.

In the current work, we proposed a new aspect extraction method for drug reviews, which is an extension of the *Aspectator* method (Bancken, Alfarone, and Davis 2014). This work was originally developed to automatically detect aspects of products on user feedback. It analyzes dependency paths in the syntactic tree of a review to find opinions expressed on candidates aspects.

The steps of this solution are shown in Figure 1. Initially, drug reviews are collected from a public repository. The reviews are divided into sentences, cleaned by removing special characters. Then tokenization and POS tagging is performed on each sentence (in our work Stanford CoreNLP tool was adopted). Spelling check is performed (in our work using the Google's spell checker). Finally, aspects are extracted by matching paths in dependency trees.

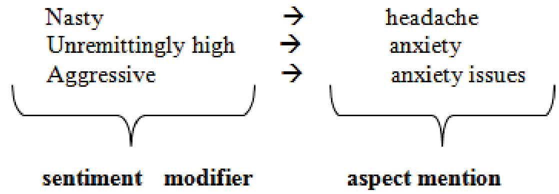

Figure 2: Opinion pair

The selection and extraction is based on pairs of words or phrases called *opinion pairs* as illustrated in Figure 2. The first term called *sentiment modifier* is the word around the aspect that expresses an opinion and the second term called *aspect mention* is the mention of an aspect.

Aspectator is capable of extracting candidate aspect terms by combining certain syntactic dependency paths. Compared to other approaches, it is not necessary to have labeled data and the algorithm also does not require domain specific knowledge (Bancken, Alfarone, and Davis 2014).

Figure 3 presents the original dependency paths (given a dependency tree) proposed by *Aspectator*. The paths "amod", "nsubj-dobj", "nsubj-xcomp" and "nsubj-cop" are represented by a pair of substantive (NN) and adjective (JJ) and the path "nsubjpassa-advmod" is represented by a pair of substantive and verb (VB). For example, the sentence *I have been having pretty bad chest pains*, it is extracted the opinion pair <bad JJ ; pains NN>by dependency path "amod". Extensions dependency paths are dependent of main paths: "compound noun" path extracts compound noun (NN) to aspect mention, "adverbial modifier" extracts a modifier (RB) to sentiment term, "simple negation" and "negation through no determiner" extract negative term related to the respective opinion pair. Therefore, the sentence cited above results the opinion pair <pretty bad JJ; chest pains NN>by main and extension paths.

The original approach focuses on extracting opinion pairs based on nouns and adjectives, except to "nsubjpass-advmod" path. In our domain of interest, Sarker (2015) comments that verbs have an important role when patients express their experiences in a review. Thus, important aspects can not be extracted by the original *Aspectator*. For instance, in the expression *reduced my pain*, the term *reduced* is a verb (a verb not related by an adverb as suggested in"nsubjpass-advmod" path) and thus the effectiveness aspect of the drug would not be identified. This paper proposes to extend the algorithm and adapt it to the medical domain in order to overcome this limitation as well as other gaps identified.

Figure 4 presents new dependency paths proposed in our work. All new dependency paths considered relations between nouns and adjective or verbs. The dependency path "amod-nsubj" relations "amod" and "nsubj". In the dependency paths "nsubj-xcomp VB-JJ" and "nsubj-xcomp VB-VB", a verb is considered as aspect mention and an adjective or a verb respectively are sentiment modifiers. All new paths are matched with Extension paths suggested in original algorithm. Each new relationship was suggested by the observation that by considering only the original dependency paths, relevant opinion pairs are ignored in the medical domain.

In our work, dependency trees were computed using the Universal Dependency Relations[1] embedded in the Stanford coreNLP(Mcdonald et al. 2013).

Dependency path	Example
Main dependency	
amod	amod I 'm talking horrible behavior PRP VBP VBG JJ NN
nsubj_dobj	nsubj dobj The pain got stale DT NN VBD JJ
nsubj_xcomp	nsubj xcomp The Ritalin made me extremely irritable DT NNP VBD PRP RB JJ
nubj_cop	nsubj cop even though my pressure is normal RB IN PRP$ NN VBZ JJ
nsubjpass_ advmod	nsubjpass advmod My anxiety was decreased tremendously PRP$ NN VBD VBN RB
Extensions dependency	
compound noun	amod compound it gave me bad chest pain PRP VBD PRP JJ NN NN
adverbial modifier	advmod had really bad headaches VBN RB JJ NNS
simple negation	neg The nausea is n't so bad DT NN VBZ RB RB JJ
Negation through "no" determiner	neg with no co-morbid ADHD IN DT JJ NN

Figure 3: Main and Extensions Dependency paths used by by Aspectator Algorithm

Experiment and discussion

In this section we describe the experiments with the proposed approach. Drug review datasets were collected from a public repository[2] regarding three clinical conditions: ADHD, Aids and Anxiety. Table 3 presents an example of review from the Anxiety dataset. In order to create a labeled corpus for evaluation, we initially split each review into sentences and manually extracted opinion pairs found. Each opinion pair was classified in one of six aspect types described in Table 1. Table 2 summarizes the datasets.

Table 2: Corpus

	ADHD	Aids	Anxiety
Drugs Name	4	4	4
Total Reviews	502	201	500
Labeled Aspect	4.680	1.895	3.985

[1]Some relations have been updated, example, Acomp to Xcomp. Full list in http://universaldependencies.org/u/dep/

[2]drugs.com

Dependency path	Dependency paths examples
amod VB	amod Some people note increased aggression DT NNS VBP VBN NN
amod_conj	conj:and amod cc Night sweats with intense anxiety and nausea NNP VBZ IN JJ NN CC NN
amod_nsubj	nsubj amod advmod The therapeutic effects are still evident DT JJ NNS VBP RB JJ
dobj NN_VB	dobj he has developed seizures PRP VBZ VBN NNS
nsubj_xcomp VB	nsubj xcomp My metabolism slowed down immensely PRP$ NN VBD IN RB
nsubj_xcomp VB_JJ	nsubj xcomp I felt very groggy PRP VBD RB JJ
nsubj_xcomp VB_VB	nsubj xcomp I mainly feel relaxed PRP RB VBP VBN
nsubj VB	nsubj neg This medicine is n't helping DT NN VBZ RB VBG
nsubj_conj	conj:and nsubj cc his behavior is uncontrollable and very irritable PRP$ NN VBZ JJ CC RB JJ

Figure 4: New Dependency Paths: Adaptation to Aspectator Algorithm to Medical Domain

Table 3: Anxiety Dataset: Review Example

Brand Name	Review	Rating
Ativan	I was battling depression along with panic attacks. They first had me on Xanax. After a while wanted to get off that but was still having panic attacks. Not since I started Ativan. Also I find it much easier to focus and get things done now.	10

All sentences were preprocessed to remove special characters and symbols, to convert capitalized terms as it may infer wrong syntactic tree parse process, to remove consecutive blank spaces between words, to perform spelling of terms and to standardize terms, as some characters can be used repeatedly in the same word. The term *meds* and *med* were replaced to *medicine*. Stopwords were not removed.

In order to evaluate the results, we computed classical metrics: precision, recall and F-measure. Precision (P) is defined by the number of automatically extracted opinion pairs that occur in the labeled dataset (i.e., relevant pairs extracted - Opa) divided by number of automatically opinion pairs extracted (Opi). Recall (R) is measured as the number of relevant opinion pairs extracted divided by the number opinion pairs present in the labeled dataset (Opn). The F-Measure (F) is the harmonic mean between precision and recall.

In order to do a comparative analysis, we also perform experiments with previous approaches:

- Zheng et al. (2014): developed an unsupervised depen-

Table 4: Precision, Recall and F-Measure to baselines, original Aspectator and new dependency paths

Dataset	Method	P	R	F
ADHD	Zheng et al. (2014)	0,459	0,561	0,505
	Samha (2016)	0,580	0,540	0,560
	Aspectator	0,816	0,514	0,631
	Method Proposed	0,78	0,665	**0,718**
Aids	Zheng et al. (2014)	0,502	0,578	0,537
	Samha (2016)	0,643	0,604	0,623
	Aspectator	0,772	0,554	0,645
	Method Proposed	0,752	0,678	**0,713**
Anxiety	Zheng et al. (2014)	0,515	0,543	0,529
	Samha (2016)	0,624	0,550	0,585
	Aspectator	0,828	0,582	0,683
	Method Proposed	0,787	0,658	**0,717**

dency analysis-based approach to extract Appraisal Expression Patterns (AEPs) from reviews regarded as a condensed representation of the syntactic relationship between aspect and sentiment words. In work, the AEP is applied to represent the syntactic relationship between aspect and sentiment words by using Shortest Dependency Path (SDP) that connects two words in the dependency graph, and it is an alternate sequence of POS tags and syntactic dependency relationships. The AEP information is incorporated into the AEP-LDA model for mining aspect and sentiment words simultaneously. We compare our dependencies patterns with relations generated in the AEP.

- Samha (2016): proposed propose a Natural Language Processing approach that undertakes Dependency Parsing, Pre-processing, Lemmatization, and part of speech tagging of natural texts in order to obtain the syntactic structure of sentences by means of a dependency relation rule. The Stanford Dependency Relations is applied to find the syntactic parsers that will allow us to map the dependencies between all words within the sentence in the form of relation (a grammatical relation holds between a head and a dependent). It's explored a set of syntactic rules and relations that were observed from the product dataset.
- Original Aspectator.

Table 4 presents the obtained results. The original *Aspectator* achieved the highest precision for all datasets, but lower recall levels. The other baselines were not competitive due to the low values of precision in turn. The proposed method obtained the best trade-off between precision and recall, i.e., the best results in terms of F-Measure.

Table 5: Recall to each aspect type

Dataset	Overall	Effectiveness	ADR	Dosage	Condition	Cost
ADHD	0,796	0,696	0,629	0,708	0,713	0,615
Aids	0,879	0,796	0,651	0,511	0,582	0,60
Anxiety	0,780	0,764	0,725	0,601	0,777	0,778

Table 5 presents the recall relative to each aspect type: aspects types retrieved by our method that occur in the labeled dataset divided by number of each type present in the labeled data set. The aspect Cost retrieved the lowest frequency where given the sum in all bases occurred only 27 pairs, users usually do not discuss about price in drug reviews. Regarding the aspect Dosage to ADHD and Anxiety had a higher occurrence than to AIDS, we observed that drugs belonging to the experiment-based ADHD and Anxiety are available in different dosages (capsule, tablet or liquid) which users have more options to discuss about different dosages.

We observed that aspect type occurrence is quite dependent of disease type. For certain aspects like condition, the recall is lower, which could be explained by the difficulty in handling some domain dependent terms. We also note that there is a number of opinions expressed about the efficacy or side effects of medications using quantitative values, which occurs mainly in the AIDS database. Note the review below:

- *My CD4 count was at 89 and my Viral Load up to 182,000/mL. I began taking Atripla on Nov. 10th 2015. After receiving my test results my CD4 count had jumped to 394, and my Viral Load dropped to 2,185 - after only 6 weeks of treatment. ...my CD4 count had jumped to 394, and my Viral Load dropped to 2,185*

implies positive review about the medicine while the *Aspectator* is not prepared to handle with quantitative terms.

Table 6: Precision by Dependency Path

Path	**ADHD**			**Aids**			**Anxiety**		
Original Paths	**Opa**	**Opi**	**P**	**Opa**	**Opi**	**P**	**Opa**	**Opi**	**P**
amod NN-JJ	1302	1627	**0,8**	591	787	**0,75**	1220	1493	**0,82**
nsubj-dobj NN-JJ	1	3	0,33	1	1	1	1	2	0,5
nsubj-xcomp NN-JJ	50	55	**0,91**	17	19	**0,89**	42	47	**0,89**
nsubj-cop NN-JJ	158	164	**0,96**	92	101	**0,91**	125	132	**0,95**
nsubjpass-advmod NN-VB	14	20	**0,7**	5	6	**0,83**	15	21	**0,71**
New Paths									
amod NN-VB	46	49	**0,94**	14	15	**0,93**	43	45	**0,96**
amod-conj NN-VB/JJ	46	73	0,63	24	50	0,48	48	74	0,65
amod-nsubj NN-VB/JJ	113	161	**0,7**	82	110	**0,75**	106	149	**0,71**
dobj NN-VB	813	1127	**0,72**	285	394	**0,72**	750	1040	**0,72**
nsubj-xcomp NN-VB	50	108	0,46	15	29	0,52	42	84	0,5
nsubj-xcomp VB-JJ	237	247	**0,96**	51	53	**0,96**	209	215	**0,97**
nsubj-xcomp VB-VB	19	25	**0,76**	2	3	0,67	5	8	0,63
nsubj NN-VB	364	907	0,4	137	319	0,43	342	767	0,45
nsubj-conj NN-VB	13	18	**0,72**	8	11	**0,73**	15	21	**0,71**

Table 6 presents the precision of each dependency path presented in Figures 2 and 3. As it can be seen, although most rules have good precision values, some of them are not so precise. For instance, as the dependency path "nsubj-doj" returned the lowest frequency and low precision, we consider this path not relevant to our domain. The dependency path "nsubj NN-VB" resulted high frequency and low precision, we observed that is necessary filtering the occurrence of some verbs types. For example, verbs "is", "do", "went", "does", "been" are frequent extracted as *sentiment modifier*.

The path "dobj NN-VB" reached high precision and solves the expression *reduced my pain* previously mentioned. The path "nsubj-xcomp VB-VB - VB-JJ" is able to extract verbs as *aspect mention*. Drugs reviews have high frequency of sentences in first person, as *I feel extremely anxious*, resulting opinion pair <extremely anxious; feel>. In addition, we considered the new dependency paths proposed relevant to drugs reviews domain.

Conclusion and further work

In this work, we extend the algorithm *Aspectator* suggesting new dependency paths to extract relevant *opinion pairs* to medical domain. We tested each new path in three datasets of drugs reviews. The proposed solution achieved very competitive results compared to baseline methods (the highest values of F-Measure were observed for all datasets). We highlight that the proposed solution can be easily adapted to other languages since it does not require labeled data.

As further work, we aim to investigate methods to automatically classify aspects types from unstructured text of drugs reviews. Other datasets can be considered in the future covering other medicines and diseases. We also aim to explore supervised machine learning, hybrid approaches and lexical resources in order to achieve improvements in results. Finally review analysis of comparative sentences can be performed in order to consider citations between drugs.

References

Asghar, M. Z.; Khan, A.; Ahmad, S.; and Ahmad, B. 2013. Subjectivity lexicon construction for mining drug reviews. *Science International* 26:145–149.

Asghar, M. Z.; Ahmad, S.; Qasim, M.; Zahra, S. R.; and Kundi, F. M. 2016. Sentihealth: creating health-related sentiment lexicon using hybrid approach. *SpringerPlus* 5(1):1139.

Bancken, W.; Alfarone, D.; and Davis, J. 2014. Automatically detecting and rating product aspects from textual customer reviews. *International Workshop on Interactions between Data Mining and Natural Language Processing*.

Cheng, V. C.; Leung, C. H. C.; Liu, J.; and Milani, A. 2014. Probabilistic aspect mining model for drug reviews. *IEEE Trans. Knowl. Data Eng.* 26(8):2002–2013.

Denecke, K., and Deng, Y. 2015. Sentiment analysis in medical settings: New opportunities and challenges. *Artificial Intelligence in Medicine* 64(1):17–27.

Egger, D.; Uzdilli, F.; and Cielebak, M. 2015. Adverse drug reaction detection using an adapted sentiment classifier. *Social Media Mining Shared Task Workshop at the Pacific Symposium on Biocomputing*.

Gosal, G. P. S. 2015. Opinion mining and sentiment analysis of online drug reviews as a pharmacovigilance technique. *International Journal on Recent and Innovation Trends in Computing and Communication* 3:4920–4925.

Harpaz, R.; Callahan, A.; Tamang, S.; Low, Y.; Odgers, D.; Finlayson, S.; Jung, K.; LePendu, P.; and Shah, N. 2014. Text mining for adverse drug events: the promise, challenges, and state of the art. *Drug safety: An Intern Jour. of medical Toxicology and Drug Experience* 37(10):777–790.

Jonnagaddala, J.; Jue, T. R.; and Dai, H.-J. 2016. Binary classification of twitter posts for adverse drug reactions. *Social Media Mining Shared Task Workshop at the Pacific Symposium on Biocomputing*.

Liu, B. 2012. *Sentiment Analysis and Opinion Mining*. Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers.

Mcdonald, R.; Nivre, J.; Quirmbach-brundage, Y.; Goldberg, Y.; Das, D.; Ganchev, K.; Hall, K.; Petrov, S.; Zhang, H.; Tckstrm, O.; Bedini, C.; Bertomeu, N.; and Lee, C. J. 2013. Universal dependency annotation for multilingual parsing. In *In Proc. of ACL '13*.

Na, J.-C., and Kyaing, W. Y. M. 2015. Sentiment analysis of user-generated content on drug review websites. *Journal of Information Science Theory and Practice* 1(1).

Na, J.-C.; Kyaing, W. Y. M.; Khoo, C. S. G.; Foo, S.; Chang, Y.-K.; and Theng, Y.-L. 2012. Sentiment classification of drug reviews using a rule-based linguistic approach. *Lecture Notes in Computer Science* 7634:189–198.

Nikfarjam, A.; Sarker, A.; O'Connor, K.; Ginn, R.; and Gonzalez, G. 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* 22(3):671–681.

Noferesti, S., and Shamsfard, M. 2015. Resource construction and evaluation for indirect opinion mining of drug reviews. *PLoS ONE* 10.

Pang, B., and Lee, L. 2008. Opinion mining and sentiment analysis. *Found. Trends Inf. Retr.* 2(1-2):1–135.

Patki, A., and Gonzalez, G. 2014. Mining adverse drug reaction signals from social media: Going beyond extraction. *22nd Annual International Conference on Intelligent Systems for Molecular Biology*.

Samha, A. K. 2016. Aspect-based opinion mining using dependency relations. *International Journal of Computer Science Trends and Technology (IJCST)* 4.

Sampathkumar, H.; Chen, X.-w.; and Luo, B. 2014. Mining adverse drug reactions from online healthcare forums using hidden markov model. *BMC Medical Informatics and Decision Making* 14(1):91.

Sarker, A., and Gonzalez, G. 2015. Portable automatic text classification for adverse drug reaction detection via multi-corpus training. *J. of Biomedical Informatics* 53:196 – 207.

Sarker, A.; Ginn, R.; Nikfarjam, A.; OConnor, K.; Smith, K.; Jayaraman, S.; Upadhaya, T.; and Gonzalez, G. 2015. Utilizing social media data for pharmacovigilance: A review. *Journal of Biomedical Informatics* 54:202–212.

Sharif, H.; Zaffar, F.; Abbasi, A.; and Abbasi, A. 2014. Detecting adverse drug reactions using a sentiment classification framework. *The sixth ASE international conference on social computing (SocialCom)*.

Veloso, A., and Jr., W. M. 2007. Efficient on-demand opinion mining. In *XXII Simpósio Brasileiro de Banco de Dados*, 332–346.

Yang, H., and Yang, C. C. 2015. Mining a weighted heterogeneous network extracted from healthcare-specific social media for identifying interactions between drugs. In *ICDM Workshops*, 196–203. IEEE Computer Society.

Zheng, X.; Lin, Z.; Wang, X.; Lin, K.-J.; and Song, M. 2014. Incorporating appraisal expression patterns into topic modeling for aspect and sentiment word identification. *Knowl.-Based Syst.* 61:29–47.