

Treatment Side Effect Prediction from Online User-generated Content

Nguyen Van Hoang Kazunari Sugiyama Min-Yen Kan Kishalay Halder

School of Computing
National University of Singapore
hoang_van_nguyen@u.nus.edu, {sugiyama, kanmy, kishalay}@comp.nus.edu.sg

Abstract

With Health 2.0, patients and caregivers increasingly seek information regarding possible drug side effects during their medical treatments in online health communities. These are helpful platforms for non-professional medical opinions, yet pose risk of being unreliable in quality and insufficient in quantity to cover the wide range of potential drug reactions. Existing approaches which analyze such user-generated content in online forums heavily rely on feature engineering of both documents and users, and often overlook the relationships between posts within a common discussion thread. Inspired by recent advancements, we propose a neural architecture that models the textual content of user-generated documents and user experiences in online communities to predict side effects during treatment. Experimental results show that our proposed architecture outperforms baseline models.

1 Introduction

Seeking medical opinions from online health communities has become commonplace: 71% of age 18–29 (equivalent to 59% of all U.S. adults) reported consulting online health opinion (Fox and Duggan, 2013). These opinions come from an estimated twenty to one hundred thousand health-related websites (Diaz et al., 2002), inclusive of online health communities that network patients with each other to provide information and social support (Johnston et al., 2013). Platforms such as HealthBoards¹ and MedHelp² feature users reporting their own health experiences, inclusive of their self-reviewed drugs and medical treatments. Hence, they are valuable sources for researchers (Leyens et al., 2017; Martin-Sanchez and Verspoor, 2014).

Although readers use these platforms to get valuable information about potential drug reactions during treatment, this is not without potentially serious problems. There is lexical variation:

users do refer side effects differently: “*dizziness*” can be expressed as “*giddiness*” or “*my head is spinning*”. More concern is that discussions rarely cover all possible prescribed drugs and their side effects during a treatment, and some topics refer to a condition without mentioning any particular drug. Relying on such information could lead to adverse reactions.

It is important to note that a tool that looks up mentioned drugs’ side effects from a static database would not return answers with sufficient coverage. There are also common concerns regarding credibility of user-generated contents – (Impicciatore et al., 1997) have shown that online health information is of variable quality and approached with caution.

Having these caveats in mind though, experienced users can provide valuable expertise. For instance, while reporting expected side effects for a specific treatment, patients with long-term use of certain drugs can be valuable authorities. *E.g.:*

While my experience of 10 years is with Paxil, I expect that Zoloft will be the same. You should definitely feel better within 2 weeks. One way I found to make it easier to sleep was to get lots of exercise. Walk or run or whatever to burn off that anxiety. – User 3690.

This is an answer to a thread asking for expected side effects for depression treatment with Zoloft. User 3690’s history of actively discussing about other anti-depressants such as Lexapro and Xanax gives insights in predicting potential drug reactions during the treatment of depression. Table 1 shows that Zoloft (mentioned in the thread) shares many common side effects with the other two anti-depressants: “*changed behavior*,” “*dry mouth*,” and “*sleepiness or unusual drowsiness*.”

A method that could differentiate trustworthy user-generated content would be valuable, allowing us to macroscopically harness a large amount of online information that would pave the way to many critical tasks such as digital pharmacovigilance (Salathé, 2016) and disease monitoring (St Louis and Zorlu, 2012). Even on the micro-

¹<https://www.healthboards.com/>

²<https://medhelp.org/>

Drugs	Side effects
Lexapro	chills, constipation, cough, decreased appetite, decreased sexual desire, diarrhea, dry mouth , joint pain, muscle ache, tingling feeling, sleepiness or unusual drowsiness , unusual dream, sweating , ...
Xanax	abdominal or stomach pain, muscle weakness, changed behavior , chills, cough, decreased appetite, decreased urine, diarrhea , difficult bowel movement, cough, dry mouth , tingling feeling, sleepiness or unusual drowsiness , slurred speech, sweating, yellow eye...
Zoloft	changed behavior , decreased sexual desire, diarrhea, dry mouth , heartburn, sleepiness or unusual drowsiness, sweating ...

Table 1: Side effects of anti-depressants.

scopic level of individual posts, such a tool offers users’ suggestions for drug reactions and improves the quality of user-generated content.

We address this need in our work. We build a neural architecture that models each post’s textual content and its author’s experience to predict expected side effects during treatments. Crucially, our supervised neural approach *jointly* learns posts’ content and users’ experience level within a thread. A key observation we make is that users can be grouped into clusters that share the same expertise or interest in certain drugs, possibly due to their common treatment or medical history. We leverage this expertise by embedding it into a low dimensional vector learned by the model, and subsequently predict side effects that are unmentioned in the discussion. We believe that our model represents trustworthiness more robustly when compared with representations such as a single weights (Li et al., 2016) and traditional drug side effect extraction (Aramaki et al., 2010). Furthermore, inspired by (Halder et al., 2018), we train a cluster-sensitive attention mechanism that allows our model to emphasize varied parts of the post. We also follow general definition of truth discovery and let the model learn a credibility score that is unique to every user and reflective of her trustworthiness. Our experimental results show that integrating the above components outperforms baseline text classification models.

The contributions of our work are summarized as follows:

- We propose a neural network architecture that can capture user expertise, user credibility, individual post’s and overall thread’s semantic content.
- We formulate the task of side effect prediction during treatment as supervised multi-label classification and apply our proposed method to the task of side effect prediction during treatment.
- We record and analyze the performance of our proposed model through a set of progressively designed experiments. Additionally, we compare the obtained results with traditional text encoding algorithms.

2 Related Work

Our approach learns the representation of posts, threads and users, and then integrates them to apply to the task of drug side effect prediction during treatment. We thus review works on the representation of fundamental objects in online communities, and the discovery of drug side effects.

2.1 Modeling Objects in Online Communities

Post content modeling. In statement credibility prediction, linguistic features of a post are strong indicators for reliability. Stylistic features – *i.e.*, the number of strong/weak modals, conditionals or negations – and affective features – *i.e.*, words that depict an author’s attitude and emotion – are adopted to represent a post’s content (Mukherjee et al., 2014). Such feature engineering requires a great amount of correlation analysis when applied to a novel problem or dataset. Linguistic features also often fail to fully capture document content, as most do not account for distinctive words in exchange of scalability. Its counter parts, bag of words and per-vocabulary features loosely capture textual content but disregard semantics and suffer scalability with sparsity issues. To address this, state-of-the-art architectures feature complex modeling to model subtle dependencies and rely on word embeddings to address scalability issues, achieving robust results in text classification (Kim, 2014), neural machine translation (Luong et al., 2016), among others.

Inspired by the success of their approaches, we adopt the recurrent neural network architecture (RNN) for post content modeling. Coupled with an attention mechanism, our approach adaptively weights the importance of parts in each post (Luong et al., 2015).

Thread content modeling. Most research working on thread-level modeling usually obtain thread content representation by aggregating each content of its posts (Yang et al., 2014). However, we hypothesize that each post has different contribution to thread content and should be variously weighted to reflect specific factors, such as its author’s level of credibility.

User ID	Post	Drug mentioned	Aggregated side effects
3690	While my experience of 10 years is with Paxil, I expect that Zoloft will be the same. You should definitely feel better within 2 weeks. One way I found to make it easier to sleep was to get lots of exercise. Walk or run or whatever to burn off that anxiety.	Zoloft	changed behavior, decreased sexual desire, diarrhea, dry mouth, heartburn, sleepiness or unusual drowsiness,...
26521	I've heard of people going "cold turkey" and having withdrawal at 6 months! Please, get in contact with a doctor ASAP! "common symptoms include dizziness, electric shock-like sensations, sweating, nausea, insomnia, tremor, confusion, nightmares and vertigo"		

Table 2: A sample thread, including its list of post–user pairs, mentioned drugs, and side effects.

User modeling. Statement credibility prediction often represents users by a single scalar that indicates their trustworthiness. The intuition is that users who provide trustworthy information frequently will be assigned high reliability scores (Li et al., 2017). Such representation is effective yet insufficient. Recent work have shown that encoding users into high-dimensional embeddings can improve system performance (Yu et al., 2016), which we have adopted in our model.

2.2 Side Effect Discovery

Most drug reaction discovery methods focus on extracting mentioned side effects. A common technique is to apply Named Entity Recognition (NER) and Relation Extraction (RE) systems in a supervised manner. (Sampathkumar et al., 2014) demonstrates its effectiveness in detecting drugs and side effects that appear in a target document (in-context), and predicting if they are related.

However, in our side effect prediction during treatment, our model is required to cover potentially encountered reactions, many of which are not explicitly mentioned in the given post (out-of-context). Hence, we do not identify our task with traditional task of adverse drug side effect extraction (Leaman et al., 2010). Our approach overcomes the limitations of the existing works by modeling user experience, and credibility during post and thread encoding, then subsequently predicting both in- and out-of-context side effects.

3 Preliminaries

Basic Terminologies. To ensure a consistent representation, let us first define some terminology:

- A drug d has a set of side effects, $S_d = \{s_1, s_2, \dots, s_k\}$
- A post p is the most basic document, containing a sequence of sentences. It is written by a user u , and belongs to a thread t .
- A user u is a member of an online community. She participates in certain threads, *i.e.*,

$T_u = \{t_1, t_2, \dots, t_l\}$ by writing at least one post in each thread. We use the terms *user* and *author*, as well as *user experience* and *user expertise* interchangeably.

- A thread t (see Table 2) is an ordered collection of post–user pairs,

$$Q_t = [(p_1, u_1), (p_2, u_2), \dots, (p_n, u_n)].$$

Every thread discusses the treatment of a particular condition and entails a list of prescribed drugs $D_t = \{d_1, d_2, \dots, d_m\}$. Hence, every thread has a list of aggregated side effects defined as $S_t = S_{d_1} \cup S_{d_2} \dots \cup S_{d_m}$, which is also the list of potential side effects experienced during the treatment.

Task Definition. *Drug side effect prediction during treatment* is the task of assigning the most relevant subset of side effects to threads discussing certain treatment, from a large collection of potential side effects. We view the drug side effect prediction problem as a multi-label classification task. In our setting, an instance of item–label is a tuple (x_t, y) where x_t is the feature vector of thread t derived from its list of post–user pairs Q_t and y is the side effect label vector *i.e.*, $y \in \{0, 1\}^S$, where S is the number of possible side effect labels. Given training instances, we train our classifier to predict the list of treatment side effects in unseen threads.

Formal Hypothesis. Given a thread t with Q_t , we hypothesize that considering the credibility and experience of user $u \in (p, u) \in Q_t$ improves the quality of feature representation in thread t , resulting in better treatment side effect prediction.

4 Proposed Method

Figure 1 shows the detailed network architecture of our model. It has several components which we shall detail sequentially. Ablation of certain components will serve as baseline systems for comparative evaluation later.

User Expertise Representation (UE): We embed each user $u \in U$ as a vector v_u so that the vector captures user u 's experience with certain drugs.

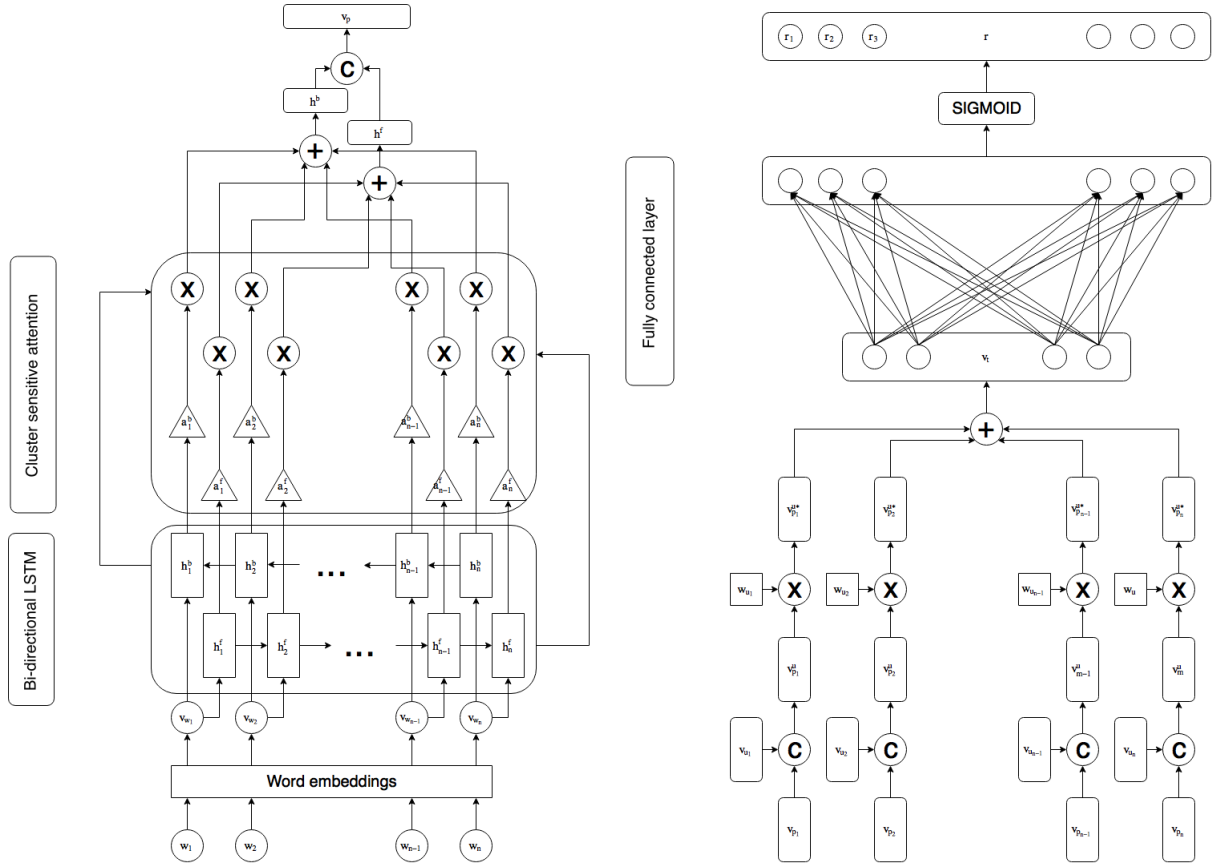


Figure 1: Full model architecture: (l, 1st half) word embeddings and cluster sensitive attention, (r, 2nd half) thread representation and multi label side effect prediction.

As each user u participates in the threads T_u , entailing a list of experienced drugs, we derive user drug experience vector $v_u^* \in \mathbb{R}^{|D|}$ where D is the set of all possible drugs and $v_{u_i}^* = n_{u_i}$ where user u has mentioned i^{th} drug in n_{u_i} threads. We obtain a user drug experience matrix $M^* \in \mathbb{R}^{|U| \times |D|}$ where j^{th} row of M^* denotes user drug experience vector of j^{th} user $u_j \in U$. Since the average number of drugs experienced per user is much fewer than the total number of drugs (see Table 3), M^* suffers from data sparsity and limited scalability. Without dimensionality reduction, the model learns at least $|D|$ parameters for every user, amounting to $|D| \times |U|$ when aggregated for all users. Data sparsity leads to a large number of insufficiently tuned parameters, which significantly increases training time, storage, and reduces the system’s robustness.

We apply Principal Component Analysis (PCA) (Jolliffe, 1986) to M^* obtained from training set. Figure 2 shows percentage of variance explained versus number of included principal components (PCs) to determine the number of PCs g . Since our PCA plots do not show added explanation percentage beyond 50 components, we use $g = 50$ com-

ponents, reducing our original $M^* \in \mathbb{R}^{|U| \times |D|}$ to user expertise matrix $M \in \mathbb{R}^{|U| \times g}$.

User Clustering: To model per-user expertise, in a naïve setting, we would train $\approx |U| \times g$ parameters. Given limited data, this is infeasible as it faces sparsity issues. We make a second, key assumption that our set of users U can be grouped into a set of meaning clusters C of size k where $k \ll |U|$. Users within a cluster would have experience with similar drugs, and hence representable using a single vector, reducing the number of learned parameters to $k \times g$.

We apply K -means clustering algorithm (MacQueen, 1967) to cluster the users into k groups. To determine the number of clusters k , we analyze the total distance to the nearest centroid versus the number of potential clusters in set C – as in Figure 3, where $D(C)$ is defined as follows:

$$D(C) = \frac{\sum_{c \in C} \sum_{u \in c} \text{dist}(v_c, v_u)}{\text{argmax} D(C)}, \quad (1)$$

where $\text{argmax} D(C)$ is the maximum total distance obtained when $|C| = 1$.

Since clustering does not gain significant reduction in total distance beyond 100 clusters, we sort




Figure 2: Principal component analysis.




Figure 3: K -means analysis.

each user to a cluster $c \in C$ where $|C| = k = 100$. For each user, we consider the vector of her assigned cluster’s centroid to be her expertise vector.

Post Content Encoding: The network takes the content of a thread t as input, which is a list of post–user pairs Q_t . Post p_i of pair $(p_i, u_i) \in Q_t$ consists of a sequence of words (w_1, \dots, w_n) . We seek to represent a post p_i as vector \mathbf{v}_p that effectively captures its semantics. We embed each word into a low dimensional vector and transform the post into a sequence of word vectors $\{\mathbf{v}_{w_1}, \mathbf{v}_{w_2}, \dots, \mathbf{v}_{w_n}\}$. Each word vector is initialized using Google’s pre-trained word2vec (Mikolov et al., 2013). Additionally, while each out-of-vocabulary word vector is initialized randomly, we keep it tunable during training to capture domain-specific meanings. Such model adaptation is necessary, as the model needs to learn the embeddings for the drug names, most of which are not included in the pre-trained embeddings but are critical to predict the side effects.

We employ Long-Short Term Memory (LSTM) (Hochreiter and Schmidhuber, 1997) to encode the textual content. A bi-directional LSTM encodes the word vector sequence, outputting two sequences of hidden states: a forward sequence, $H^f = \mathbf{h}_1^f, \mathbf{h}_2^f, \dots, \mathbf{h}_n^f$ that starts from the beginning of the text; and a backward sequence, $H^b = \mathbf{h}_1^b, \mathbf{h}_2^b, \dots, \mathbf{h}_n^b$ that starts from the end of the text. For many sequence encoding tasks, knowing both past (left) and future (right) contexts has proven to be beneficial (Dyer et al., 2015). The states \mathbf{h}_i^f and \mathbf{h}_j^b in the forward and backward sequences are computed as follows:

$$\mathbf{h}_i^f = LSTM(\mathbf{h}_{i-1}^f, \mathbf{w}^i), \quad \mathbf{h}_j^b = LSTM(\mathbf{h}_{j+1}^b, \mathbf{w}^j),$$

where $\mathbf{h}_i^f, \mathbf{h}_j^b \in \mathbb{R}^e$, and e are the number of encoder units.

# Users	# Threads	Avg. # of words per post
14,388	99,682	73.65
Avg. # of posts per thread		Avg. # of threads per user
8.16		26.21
# Side effects (SE)	Avg. # of SEs per thread	
1,500	90.47	
# Drugs	Avg. # of drugs per user	
1869	19.72	

Table 3: Dataset statistics.

Cluster-sensitive Attention (CA): Inspired by (Halder et al., 2018), we initialize an attention vector, $\mathbf{v}_{a_i} \in \mathbb{R}^e$ for each cluster c_i . Given a forward sequence $H^f = \mathbf{h}_1^f, \mathbf{h}_2^f, \dots, \mathbf{h}_n^f$ and backward sequence $H^b = \mathbf{h}_1^b, \mathbf{h}_2^b, \dots, \mathbf{h}_n^b$ of hidden post states p written by user u belonging to cluster c_i , the corresponding w_j weights each hidden state \mathbf{h}_j^f and \mathbf{h}_j^b of both sequences based on their similarity with the attention vector are:

$$w_{a_j} = \frac{\exp(\mathbf{v}_{a_i} \mathbf{h}_j)}{\sum_{l=1}^n \exp(\mathbf{v}_{a_i} \mathbf{h}_l)}. \quad (2)$$

The intuition behind Equation (2), inspired by (Luong et al., 2015), is that hidden states which are similar to the attention vector \mathbf{v}_{a_i} should be paid more attention to; hence are weighted higher during document encoding. \mathbf{v}_{a_i} is adjusted during training to capture hidden states that are significant in forming the final post representation. w_{a_j} is then used to compute forward and backward weighted feature vectors:

$$\mathbf{h}^f = \sum_j^n w_{a_j} \mathbf{h}_j^f, \quad \mathbf{h}^b = \sum_j^n w_{a_j} \mathbf{h}_j^b. \quad (3)$$

We concatenate the forward and backward vectors to obtain a single vector, following previous bi-directional RNN practice (Ma and Hovy, 2016).

Thread Content Encoding with Credibility Weights (CW): For every post–user pair (p_i, u_i) of thread t , we first compute post p_i feature vector v_{p_i} . It is then concatenated with user u_i ’s expertise vector v_{u_i} to form post–user complex vector $v_{n_i}^p$. This user-post complex is weighted by a user credibility score w_{u_i} , which is initially randomized and updated while training, to obtain final post–user pair representation $v_{n_i}^{p*}$. This follows the general intuition from the truth discovery literature that users providing high quality answers should assign higher credibility scores, and answers from credible users are more significant. Thus, the thread content representation can be defined as the weighted sum of each post–user complex vector:

$$v_t = \sum_{i=1}^n v_{n_i}^{p*} = \sum_{i=1}^n w_{u_i} v_{n_i}^p. \quad (4)$$

Multi-label Prediction: We feed the thread content representation v_t through a fully connected layer which outputs can be computed as:

$$s_t = W \tanh(v_t) + b, \quad (5)$$

where W and b are weights and biases of the layer. The output vector $s_t \in \mathbb{R}^{|S|}$ is finally passed through a sigmoid activation function, and trained using cross-entropy loss as defined by L :

$$L = \frac{1}{T} \sum_{t=1}^T (y_t \cdot \log(\sigma(s_t)) + (1 - y_t) \cdot \log(1 - \sigma(s_t))) + \lambda \sum_u v_u^2 \quad (6)$$

We adopt regularization that penalizes the training loss with the user experience matrix’s L_2 norm by a factor of $\lambda = 0.0065$, obtained via hypertuning. The loss function is differentiable, thus trainable with Adam (Kingma and Ba, 2015). During our gradient-based learning, user credibility score w_{u_i} of user u_i can be updated by calculating $\frac{\partial L}{\partial w_{u_i}}$ by back-propagation:

$$\begin{aligned} \frac{\partial L}{\partial w_{u_i}} &= \frac{\partial L}{\partial s_t} \frac{\partial s_t}{\partial v_t} \frac{\partial v_t}{\partial w_{u_i}} \\ &= \frac{\partial L}{\partial s_t} W (1 - \tanh^2(v_t)) v_{n_i}^p \end{aligned} \quad (7)$$

5 Experiments

We conduct experiments to validate the effectiveness of our proposed model. In specific, (1) we want to compare our architecture with text encoding baselines, (2) highlight performance improvements incrementally, and (3) evaluate and analyze

the obtained results, both at the macroscopic and microscopic levels.

5.1 Baselines

As a competitive baseline from prior work, CNN-KIM (Kim, 2014) constructs a document matrix that incorporates word embeddings, then applies a convolution filter to obtain feature maps. These feature maps are passed through a max-pooling filter to construct a document representation. During prediction, the representation is fed through a fully connected layer. We replace the final softmax layer of the author’s model with sigmoid to make it work in a multi-label prediction setting.

The following baselines are used to perform an ablation study of our model.

- **RNN:** We implement a bi-directional LSTM baseline, which is equivalent to our proposed method without CA, UE and CW.
- **Weighted Post Encoder (WPE):** We construct thread representation by summing each of its post–user complex vector weighted by user credibility. This is equivalent to our proposed methodology without CA and UE.
- **Weighted Post Encoder with User Expertise (WPEU):** We concatenate user expertise with post vector to create post–user complex vector. This is equivalent to our proposed method without CA.

5.2 Dataset

We conduct our experiments on the same dataset as (Mukherjee et al., 2014) including 15,000 users and 2.8 million posts extracted from 620,510 HealthBoards¹ threads.

Ground truth possible side effects experienced during treatment are defined as the side effects of drugs mentioned in the discussion. As annotating such amount of posts is expensive, drug side effects are extracted from Mayo Clinic’s Drugs and Supplements portal³ and are used as surrogates for potential reactions of treatments.

5.3 Experimental Settings

We applied a standard natural language preprocessing — Snowball stemming (Porter, 1980) and stop-word elimination — before representation modeling. From the original dataset, we only extract threads that are annotated with drugs and their side effects, along with the lists of contained posts and corresponding users. Table 3 shows the dataset statistics. We divide our data into 10

³<https://www.mayoclinic.org/drugs-supplements>

System	Components			Experiment 1			Experiment 2		
	CW	UE	CA	Pre.	Rec.	F_1	Pre.	Rec.	F_1
1. CNN-KIM				0.818	0.677	0.751	0.813	0.503	0.614
2. RNN				0.810	0.657	0.735	0.808	0.484	0.599
3. WPE	✓			0.873	0.678	0.773	0.859	0.507	0.638
4. WPEU	✓	✓		0.865	0.705	0.781	0.819	0.537	0.643
5. Our model	✓	✓	✓	0.844	0.730	0.793	0.788	0.573	0.659

Table 4: Experimental results with both actual (Experiment 1) and Strict (Experiment 2) settings. In the Component columns, “CW”, “UE”, “CA” denote “Credibility Weights”, “User Expertise” and “Cluster Attention module components”, respectively.

folds to perform cross-validation (8,1,1 folds for training, validation, and testing respectively). We perform PCA and K -means clustering on training set, using scikitlearn’s built-in modules (Pedregosa et al., 2011), with $g = 50$ principal components and $k = 100$ clusters.

For CNN-KIM, we experiment with filters with varying window sizes from 2 to 5, and set the number of feature maps for each filter to 256 and dropout to 0.5. For our proposed model and baseline models using the RNN architecture, when performing post content encoding, we set the number of units in the LSTM cell to 128. Dropout rates of 0.2 and 0.5 are used in our LSTM cells and FC layers, respectively. Cluster attention vectors and user credibility values are initialized with values ranging from -1.0 to 1.0. For each user u , we initialize her expertise vector with the value of v_u obtained in Section 4 and allow training to fine-tune. All models are trained using Tensorflow⁴ library.

We conducted two separate experiments:

- **Experiment 1:** We keep the text as-is. Any mentioned drugs are retained inside the thread.
- **Experiment 2:** We remove all mentions of any drug in our drug list. This is a more aggressive experiment which asks the model to predict the treatment’s side effects without any mention of the experienced drugs.

6 Results and Evaluation

Table 4 shows the precision, recall, and F_1 obtained by our method and the four baselines.

Macroscopic Analysis: Firstly, all of the three models that apply credibility weighting (CW) – WPE, WPEU and our model – outperform both RNN and CNN baselines in both experiments. Specifically, weighting each post by its author credibility improves the performance of naive post encoder by 6.32%, 2.15% and 3.86% on precision,

recall and F_1 respectively for Experiment 1. Results for Experiment 2 are similar. This demonstrates the effectiveness of accounting for author credibility when encoding thread content, improving side effect prediction.

Improvements by incorporating user experience (UE) are less pronounced. In Experiment 1, adding UE (WPEU vs. WPE) improves recall by 2.65% and 0.8% in F_1 . Again, the stricter Experiment 2 shows similar performance trends. On a macro scale, these statistics indicate that our model successfully learns to include more side effects in its prediction, where many are relevant to the ground truth. This is consistent with our hypothesis that considering author experience of each post is effective in predicting out-of-context side effects.

Applying cluster-sensitive attention (CA) in combining RNN’s hidden states also improves the performance. In Experiment 1, we observe that adding CA (our model vs. WPEU) also improves recall and F_1 , where again, Experiment 2 demonstrates similar but slightly more pronounced performance changes. These indicate that the attention mechanism is more effective when the drugs are present since the drug names in our documents are the phrases that receive greater emphasis.

As settings in Experiment 1 start with more information compared with those in Experiment 2, the task is easier and thus performance is improved (12.7% to 14.15% in F_1). The margin for improvement for Experiment 2 is larger, which explains why absolute score improvements are larger in Experiment 2. When measuring relative improvement, the gains are comparable.

Generally, according to the macroscopic analysis of results in Table 4, we conclude that all of the three components in our proposed architecture, namely, CW, UE, and CA have a positive impact on the overall performance of the model. We observe consistent improvements in F_1 after adding each component is consistent with our stated hypotheses, in both experimental settings.

⁴<https://www.tensorflow.org/>

User ID	Posts	Output side effects					
		CNN-KIM	RNN	WPE	WPEU	Our model	Ground truth
24296 (cred- ibility: 0.11)	[...] little red rashes all over my body that resembled vasculitis. [...] I was diagnosed and treated with the "standard treatment" twice, to not much effect), a very stiff neck, really bad brain fog and confusion . [...]	diarrhea, skin rash	skin rash	headache, diarrhea, skin rash	headache, diarrhea, unusual tiredness and weakness, dizziness, sleepiness, fever, nausea, bad breath	headache, diarrhea, unusual tiredness and weakness, dizziness, sleepiness, fever, nausea, heartburn, belching, indigestion, acid stomach, difficult bowel movement, bad breath, bone joint pain	headache, diarrhea, unusual tiredness and weakness, dizziness, fever, nausea, loss appetite, chills, heartburn, belching, indigestion, acid stomach, confusion, skin rash, weight loss, difficult bowel movement, shakiness
1537 (cred- ibility: 0.32)	[...] now last month my symptoms including joint pains, twitching and tremors and bug crawling under my scalp sensations reappeared [...]						
5232 (cred- ibility: 0.36)	[...] I don't know about cysts in the brain per se [...]						
16248 (cred- ibility: 0.21)	[...] I've been growing increasingly sensitive to more foods over the last year [...] How do you know that you had damage to your intestines from Lyme? [...] I'm curious because I am in the process of getting a Lyme work up and my intestines are messed up , but all GI tests came back negative.						

Table 5: A sample thread in the test set, mentioning drugs *Flagyl*, *Tinidazole*, *Plaquenil*, and *Vitamins*.

User ID	Experienced drugs	Top common experienced side effects
24296	rifampin, vitamin , clarithromycin, aciphex, a zithromax, plaquenil, flagyl , minocycline, levaquin, tetracycline, tinidazole , advil	diarrhea, bad breath, headache, heartburn, unusual tiredness and weakness, nausea, fever
1537	vitamin , rocephin, hydroquinone, plaquenil, flagyl , minocycline, levaquin, tinidazole	diarrhea, skin rash, headache, dizziness, heartburn, bad breath, sleepiness
5232	doxycycline, prozac, vitamin , norvasc, tylenol, flagyl , questran, biotin, cefuroxime, plaquenil	bad breath, diarrhea, nausea, dizziness, unusual tiredness and weakness
16248	celexa, prilosec, vitamin , rocephin, klonopin, nexium, fumarate, elidel, citrate, prozac	diarrhea, sneezing, nausea, excessive gas, body pain, loss voice, heart burn

Table 6: Experienced drugs and common side effects among users.

Microscopic Analysis: We also analyze our model performance at per-sample level to check whether they are consistent with the macroscopic results. We aim to confirm three hypotheses: (1) Considering author expertise improves prediction on out-of-context side effects. (2) Considering author credibility improves the extraction of both in- and out-of-context side effects from trustworthy users' content. (3) Placing attention on different parts of the document enhances the performance of in-context side effect extraction. Tables 5 and 6 show a sample testing thread, its users' commonly experienced drugs, and its side effects.

We observe that CNN-KIM and the simple, RNN-based post encoding can capture side effects that are mentioned both directly (e.g., "skin rash") as well as indirectly (e.g., "diarrhea"), but fail to capture the remaining symptoms, many of which are out-of-context.

Considering User 1537's credibility shows performance improvements. In her posts, User 1537 indirectly refers to "headache" by mentioning "bug crawling under my scalp sensations". The calculated higher credibility score weights User 1537 experiences with "sleepiness" higher in the WPEU (CW + UE) baseline prediction, which is correct. These observations are consistent with

our hypothesis about user credibility.

User experience is effective in predicting out-of-context symptoms. In the illustrated sample training set, all of the four users have experience with similar drugs with common side effects such as "unusual tiredness and weakness", "nausea", and "fever". As "bad breath" is also a shared side effect, it is comprehensible that the model outputs "bad breath". Nonetheless, it is intuitive for the model to pick up such commonness among users and compute relevant results. These observations are consistent with our hypothesis on user experience.

Finally, the model with CA can learn different parts of the documents. Especially for User 16248's posts that mentioned digestive problems, hidden states encode phrases such as "increasingly sensitive to more foods", and "damage to your intestines" receive higher attention, resulting in the prediction of "heartburn", "belching", "indigestion", "acid stomach", and "difficult bowel movement". This functionality is consistent with our original purpose and expectation for adding attention to the post encoder architecture.

7 Conclusion

We have addressed the importance of user experience and credibility in modeling thread contents of online communities, specifically through the task of drug side effect prediction during treatment. We suggest a subset of side effects relevant to the mentioned treatment in the given discussion, taking into account the each post content and its author expertise in certain treatments. Mainstream models for online communities fail to fully capture post content semantically and user experience with previous drugs.

We model users' expertise by examining their experience with different drugs, then group users with similar experience into clusters that share a common experience vector representation. Experimental results show that our proposed thread content encoder outperforms state-of-the-art document encoders, and that our neural components play a significant role in improving task performance.

We believe that our model is adaptable to other domains. We aim to use it for downstream application in online health community such as credibility analysis and thread recommendation in the future.

References

- Eiji Aramaki, Yasuhide Miura, Masatsugu Tonoike, Tomoko Ohkuma, Hiroshi Masuichi, Kayo Waki, and Kazuhiko Ohe. 2010. Extraction of Adverse Drug Effects from Clinical Records. *MedInfo 2010*, 160:739–743.
- Joseph A. Diaz, Rebecca A. Griffith, James J. Ng, Steven E. Reinert, Peter D. Friedmann, and Anne W. Moulton. 2002. Patients' Use of the Internet for Medical Information. *Journal of General Internal Medicine*, 17(3):180–185.
- Chris Dyer, Miguel Ballesteros, Wang Ling, Austin Matthews, and Noah A Smith. 2015. Transition-Based Dependency Parsing with Stack Long Short-Term Memory. In *Proc. of the 53rd Annual Meeting of the Association for Computational Linguistics (ACL 2015)*, pages 334–343.
- Susannah Fox and Maeve Duggan. 2013. Health Online 2013. *Health*, 2013:1–55.
- Kishaloy Halder, Lahari Poddar, and Min-Yen Kan. 2018. Cold Start Thread Recommendation as Extreme Multi-label Classification. In *Proc. of the Workshop on Extreme Multilabel Classification for Social Media co-located with the Web Conference (WWW'18 Companion)*, pages 1911–1918.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Computation*, 9(8):1735–1780.
- Piero Impicciatore, Chiara Pandolfini, Nicola Casella, and Maurizio Bonati. 1997. Reliability of Health Information for the Public on the World Wide Web: Systematic Survey of Advice on Managing Fever in Children at Home. *BMJ*, 314(7098):1875.
- Allen C. Johnston, James L. Worrell, Paul M. Di Gangi, and Molly Wasko. 2013. Online Health Communities: an Assessment of the Influence of Participation on Patient Empowerment Outcomes. *Information Technology & People*, 26(2):213–235.
- Ian T Jolliffe. 1986. Principal Component Analysis and Factor Analysis. *Statistical Methods in Medical Research*, 1(1):115–128.
- Yoon Kim. 2014. Convolutional Neural Networks for Sentence Classification. In *Proc. of the 2014 Conference on Empirical Methods on Natural Language Processing (EMNLP 2014)*, pages 1746–1751.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *Proc. of the 3rd International Conference for Learning Representations (ICLR2015)*.
- Robert Leaman, Laura Wojtulewicz, Ryan Sullivan, Annie Skariah, Jian Yang, and Graciela Gonzalez. 2010. Towards internet-age pharmacovigilance: extracting adverse drug reactions from user posts to health-related social networks. In *Proc. of the 2010 Workshop on Biomedical Natural Language Processing (BioNLP 2010)*, pages 117–125.
- Lada Leyens, Matthias Reumann, Nuria Malats, and Angela Brand. 2017. Use of Big Data for Drug Development and for Public and Personal Health and Care. *Genetic Epidemiology*, 41(1):51–60.
- Yaliang Li, Nan Du, Chaochun Liu, Yusheng Xie, Wei Fan, Qi Li, Jing Gao, and Huan Sun. 2017. Reliable Medical Diagnosis from Crowdsourcing: Discover Trustworthy Answers from Non-Experts. In *Proc. of the 10th ACM International Conference on Web Search and Data Mining (WSDM 2017)*, pages 253–261.
- Yaliang Li, Jing Gao, Chuishi Meng, Qi Li, Lu Su, Bo Zhao, Wei Fan, and Jiawei Han. 2016. A survey on truth discovery. *ACM SIGKDD Explorations Newsletter*, 17(2):1–16.
- Minh-Thang Luong, Quoc V Le, Ilya Sutskever, Oriol Vinyals, and Lukasz Kaiser. 2016. Multi-task Sequence to Sequence Learning. In *Proc. of the 4th International Conference for Learning Representations (ICLR2016)*.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective Approaches to Attention-based Neural Machine Translation. In *Proc. of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP 2015)*, pages 1412–1421.

- Xuezhe Ma and Eduard Hovy. 2016. End-to-end Sequence Labeling via Bi-directional LSTM-CNNs-CRF. In *Proc. of the 54th Annual Meeting of the Association for Computational Linguistics (ACL 2016)*, pages 1064–1074.
- J. MacQueen. 1967. Some Methods for Classification and Analysis of Multivariate Observations. In *Proc. of the 5th Berkeley Symposium on Mathematical Statistics and Probability*, pages 281–297.
- F. Martin-Sanchez and K. Verspoor. 2014. Big data in medicine is driving big changes. *Yearbook of Medical Informatics*, 9(1):14–20.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *In Proc. of the Advances in Neural Information Processing Systems (NIPS 2013)*, pages 3111–3119.
- Subhabrata Mukherjee, Gerhard Weikum, and Cristian Danescu-Niculescu-Mizil. 2014. People on Drugs: Credibility of User Statements in Health Communities. In *Proc. of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'14)*, pages 65–74.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research (JMLR)*, 12(2011):2825–2830.
- M. F. Porter. 1980. An Algorithm for Suffix Stripping. *Program*, 14(3):130–137.
- Marcel Salathé. 2016. Digital Pharmacovigilance and Disease Surveillance: Combining Traditional and Big-Data Systems for Better Public Health. *The Journal of Infectious Diseases*, 214(suppl.4):S399–S403.
- Hari Prasad Sampathkumar, Xue-Wen Chen, and Bo Luo. 2014. Mining Adverse Drug Reactions from Online Healthcare Forums using Hidden Markov Model. *BMC Medical Informatics and Decision Making*, 14(1):91–108.
- Connie St Louis and Gozde Zorlu. 2012. Can Twitter Predict Disease Outbreaks? *BMJ: British Medical Journal (Online)*, 344(e2353).
- Diyi Yang, Mario Piergallini, Iris Howley, and Carolyn Rose. 2014. Forum Thread Recommendation for Massive Open Online Courses. In *Proc. of the 7th International Conference on Educational Data Mining (EDM 2014)*, pages 257–260.
- Yang Yu, Xiaojun Wan, and Xinjie Zhou. 2016. User Embedding for Scholarly Microblog Recommendation. In *Proc. of the 54th Annual Meeting of the Association for Computational Linguistics (ACL 2016)*, pages 449–453.