Multi-Task Pharmacovigilance Mining from Social Media Posts

Shaika Chowdhury
Department of Computer Science
University of Illinois at Chicago
Chicago, IL
schowd21@uic.edu

Chenwei Zhang
Department of Computer Science
University of Illinois at Chicago
Chicago, IL
czhang99@uic.edu

Philip S. Yu
Department of Computer Science
University of Illinois at Chicago
Chicago, IL
psyu@uic.edu

ABSTRACT

Social media has grown to be a crucial information source for pharmacovigilance studies where an increasing number of people post adverse reactions to medical drugs that are previously unreported. Aiming to effectively monitor various aspects of Adverse Drug Reactions (ADRs) from diversely expressed social medical posts, we propose a multi-task neural network framework that learns several tasks associated with ADR monitoring with different levels of supervisions collectively. Besides being able to correctly classify ADR posts and accurately extract ADR mentions from online posts, the proposed framework is also able to further understand reasons for which the drug is being taken, known as ‘indications’, from the given social media post. A *coverage-based* attention mechanism is adopted in our framework to help the model properly identify ‘phrasal’ ADRs and Indications that are attentive to multiple words in a post. Our framework is applicable in situations where limited parallel data for different pharmacovigilance tasks are available. We evaluate the proposed framework on real-world Twitter datasets, where the proposed model outperforms the state-of-the-art alternatives of each individual task consistently.

CCS CONCEPTS

• **Information systems** → **Content analysis and feature selection; Information extraction**; • **Computing methodologies** → **Information extraction; Multi-task learning; Neural networks**;

KEYWORDS

Multi-Task Learning, Pharmacovigilance, Adverse Drug Reaction, Attention Mechanism, Coverage, Recurrent Neural Network, Social Media

ACM Reference Format:

Shaika Chowdhury, Chenwei Zhang, and Philip S. Yu. 2018. Multi-Task Pharmacovigilance Mining from Social Media Posts. In *WWW 2018: The 2018 Web Conference, April 23–27, 2018, Lyon, France*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3178876.3186053>

1 INTRODUCTION

Various prescription drugs intended to be taken for medical treatment are released into the market. However, studies have found that many of such drugs can be counterproductive [5, 25]. The harmful

reactions or injuries that are caused by the intake of drugs are known as Adverse Drug Reactions (ADR) [31] and happens to be the fourth leading cause of death in the United States [5]. Preventive steps such as ADR monitoring and detection, which are collectively called Pharmacovigilance, are critical to ensure safety to patients’ health. Initial activities toward pharmacovigilance during clinical trials fail to reliably signal all the negative effects that could potentially be caused by a drug due to certain restrictions, calling the need for post-market ADR surveillance. Nevertheless, potentially harmful drugs remain unflagged as traditional post-market ADR monitoring methods suffer from under-reporting, incomplete data, and delays in reporting [31].

Owing to the presence of large user base on social media, a vast amount of data gets generated. Compared to clinical information retrieval from electronic health records (EHR) that pose the problem of limited access [37], the free accessibility of web data presents a lucrative source of medical data. This data has recently caught the attention of researchers for public health monitoring [17, 26]. Task-oriented crowdsourcing platforms are designed to collect patient feedbacks as introduced in [27]. A popular platform where health-related data in the form of posts/ tweets are exchanged between users is Twitter¹. The posts cover a wide range of topics concerning health such as experience getting a disease/illness, symptoms of it, the drugs that were taken and harmful reactions from them.

Social media platforms offer a robust source of health-related data that patients are found to consult to learn about other’s shared health-related experiences [16, 31, 41, 42]. The fact that this data is up-to-date and is generated by patients overcomes the weaknesses of traditional ADR surveillance techniques. Thus social media is indispensably important and could complement traditional information sources for more effective pharmacovigilance studies, as well as potentially serve as an early warning system for unknown ADRs [21]. However, users with different background knowledge or various linguistic preferences tend to generate social media posts with diverse expressions and ambiguous mentions, which pose a series of challenges for pharmacovigilance studies:

- **Diversely expressed ADRs:** In Twitter, the use of non-medical terms to craft one’s tweet is very prevalent. For example, *never sleeping* or *sleep deprived* being used to describe the ADR *insomnia*. These ‘phrasal’ ADRs framed with casual words could easily merge with other irrelevant words in the post and make the detection task more difficult.
- **Indications misidentified as ADRs:** An ‘Indication’ in a health-related post can be defined as a medical condition, disease or illness for which the drug has been prescribed to treat it. As both indications and ADRs can be referred to by

This paper is published under the Creative Commons Attribution 4.0 International (CC BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW 2018, April 23–27, 2018, Lyon, France

© 2018 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC BY 4.0 License.

ACM ISBN 978-1-4503-5639-8/18/04.

<https://doi.org/10.1145/3178876.3186053>

¹<https://twitter.com>

the same drug, indications could be easily misidentified for ADRs. Without a deep understanding of the role that each symptom plays in a post, such misidentification could be problematic.

Tweet 1: I'm starting to think my **Paroxetine** turns **panic attacks** into **fat**.
Tweet 2: Any ideas on treating **depression** naturally **Paxil** is making me **gain unwanted pounds** and **water weight**...Thanks

Figure 1: Pharmacovigilance reported in social media posts.

Figure 1 shows common tweets for Pharmacovigilance taken from PSB 2016 Social Media Shared Task Twitter dataset [32], where the token in red denotes the 'Drug' being taken, blue denotes the 'Indication', and green the 'ADR'. In Tweet 1, the drug *Paroxetine* is being taken to treat *panic attacks*, with an adverse reaction of weight gain mentioned as *fat*. So *panic attacks* is an instance of Indication while *fat* is the ADR. In Tweet 2 while the drug *Paxil* is prescribed to treat the indication *depression*, it causes ADRs of *gain unwanted pounds* and *water weight*. In both cases it is possible for the Indications 'panic attacks' and 'depression' to get mislabeled as both commonly occur as ADRs as well.

All previous works [7, 13, 21, 22] in pharmacovigilance studies have focused on solving the ADR tasks separately. As a tweet containing ADR mention can also have other medical mentions such as Indications or beneficial effects, it is important to incorporate these representations to ensure disambiguity when trying to learn each task. Lexicon-based approaches [3, 17, 39] have been commonly adopted in earlier approaches, where ADR tokens in a text fragment are looked up in a corpus containing ADR mentions. However, the occurrence of non-medical terminologies to describe ADRs in social media makes them unsuitable. While machine learning approaches such as Naive Bayes, Support Vector Machine and Maximum Entropy [4, 14, 40] require the use of hand-engineered features.

Trying to conquer these problems, we propose a multi-task framework based on the sequence learning model that jointly learns several related pharmacovigilance tasks. Mining from pharmacovigilance forms an interesting multi-task problem as complimentary pharmacovigilance tasks with supervision at different levels can be jointly learned. The data for these related tasks share semantic and syntactic similarities, and leveraging these shared representations can improve learning efficiency and prediction accuracy. The joint learning of the objective functions of the tasks can transfer what is learned from one task to other tasks and improve them. This is suitable for pharmacovigilance tasks as they contain data with occurrences of multiple medical terms, which otherwise makes identification of a certain category (i.e ADR or Indication) of medical term difficult. We incorporate these ideas by extending the basic recurrent neural network encoder-decoder such that our multi-task model shares an encoder among all the tasks and uses a different decoder for each task. Our assumption is that the shared encoder

would learn predictive representations capturing the nuances of each task and, hence, help disambiguate an 'ADR' from an 'Indication'. On the other hand, a task-specific decoder decodes the shared encoder representation to produce task-specific output. Furthermore, by incorporating multi-granular supervision through different tasks, the decoder can successfully produce output at the sentence and word level.

Recent years have seen a spike in the use of sequence-to-sequence models for difficult learning tasks like machine translation [34] and summarization [20]. A sequence-to-sequence model takes a sequence as input into an encoder, projects it to an intermediate encoded representation, which is then passed through a decoder to generate a sequence as the output. The proposed architecture shown in Figure 2 is a recurrent neural network-based encoder-decoder model augmented with attention and coverage mechanism to handle multiple tasks for pharmacovigilance. The reason for using a sequence to sequence model to model the multiple tasks is to be able to capture their shared representations with the encoder and generate outputs for each task with multiple decoders. We propose a multi-task learning framework for three pharmacovigilance tasks — ADR Classification, ADR Labeling and Indication Labeling— hypothesizing that the interactions between the tasks through joint learning would improve learning and generalization ability of each individual task. While most works on ADR classification and detection have tried to learn a single objective function, by jointly learning multiple objective functions for the tasks we introduce a novel approach for pharmacovigilance that is seen to boost the performance. Moreover, the learned features from multiple tasks help to reduce the false positives in mislabeling Indications as ADRs and vice versa. When an ADR\Indication occurs as a phrase rather than a single word, it can make the detection task more difficult. Adding *coverage* to the attention mechanism helps overcome this as it accumulates the attention from all the previous decoder time steps and facilitates in learning all previous ADR\Indication words in a phrase.

The main contributions of this work are summarized as follows:

- (1) Designed a unified machine learning framework to learn several pharmacovigilance tasks from social media posts simultaneously. To the best of our knowledge, this problem has not been studied carefully and has scope for novel research.
- (2) Adding *coverage* to attention mechanism has shown to improve detection of not only 'phrasal', but also single worded ADRs and Indications.
- (3) State-of-the-art in terms of results obtained on real-world Twitter datasets compared to contemporary work in pharmacovigilance.

2 PRELIMINARY: TASKS

With our multi-task framework, we jointly learn three pharmacovigilance tasks, where the input tweet representation is encoded through a shared encoder. Each task is modeled as a sequence classification or sequence labeling problem. Description of each task is given below.

Figure 2: Model Architecture. We use different colors for each task-specific decoder. The shade of the colored block indicates its value. For example, the second word *panic* is an Indication, so gets higher attention weight (the dark blue block), which helps in successfully labeling it as an Indication.

2.1 ADR Classification

This is a binary classification task to separate out the ADR assertive posts. The two classes are ‘ADR’ and ‘NotADR’ where ‘ADR’ label indicates a post with ADR mention in it. While ‘NotADR’ means it does not have any ADR, although can have other medical terms such as drugs and Indications.

It tries to learn a function that maps a sequence $\mathbf{x}$ to a class label $l \in R^L$ where L is the total number of classes [38] and $L = 2$.

$$\mathbf{x} \rightarrow l, l \in R^L \quad (1)$$

2.2 ADR Labeling

A sequence labeling task aiming to identify ADR in a post. The detected ADRs are tagged with the ‘ADR’ label. It tries to find the most likely sequence of tags given the input sequence $\mathbf{x}$, that is the sequence with the highest probability.

$$y' = \underset{y}{\operatorname{argmax}} P(y|\mathbf{x}) \quad (2)$$

We use two tags to annotate the tokens of the input for ADR labeling. ‘ADR’ tag corresponds to a token with ADR mention and ‘O’ tag denotes a non-ADR word.

2.3 Indication Labeling

A sequence labeling task aiming to identify Indication in a post. The detected Indications are tagged with the ‘Indication’ label. Like ADR labeling, it tries to find the tagged sequence of highest probability.

We used two tags to annotate the tokens of the input for Indication labeling. ‘IND’ tag corresponds to a token with Indication mention and ‘O’ tag denotes a non-Indication word.

3 OUR MODEL

Our proposed multi-task framework is depicted in Figure 2. It is composed mainly of three components — embedding module, encoder and decoder. The embedding module is intended to capture the meaning and semantic associations between pharmacovigilance

words. All the tasks have a common encoder so that shared representations can be generated, capturing both the ADR and Indication contextual information. Lastly, the decoders employ a combined attention and coverage mechanism to facilitate the detection of ADR and Indications of various lengths. Each component of our RNN Encoder-Decoder model is described in detail in the following subsections.

3.1 Input

Our input $\mathbf{x} = (x_1, x_2, \dots, x_T)$ is a piece of text corresponding to a social media post comprising a sequence of T words, where each x_i represents a word in the vocabulary of size V .

3.2 Word Representations

Medical terms in a tweet post can take different roles and the same medical word or phrase can mean different things in different context. For example, ‘panic attacks’ can occur both as an ADR as well as an Indication depending on the mention of the drug and other patterns. In order to capture their meanings as well as the semantic relationships and context they are being used in, we generate their word embeddings.

The character representation of each word is also generated to capture its morphological features. These character representations can help in capturing the representations for words such as ‘sleeping’ where the word embedding matrix might have an entry for just ‘sleep’. This character representation is similar to that implemented for doing Named Entity Recognition [29].

The final representation of a word is the concatenation of its word embedding and character representations to incorporate users’ diverse expressions.

Word Embedding: For each word x_t we get its corresponding low-dimensional dense vector, e_word_t , by looking up in a $V \times m$ size word-embedding matrix, where m is the dimensions of the word-embedding.

Character Representation: Each word in a sentence can be denoted as $x_t = (c_1, c_2, \dots)$ such that $c_t \in R^G$, where G is the vocabulary size of all the characters. Similar to word embedding, we first get the character embedding vector for each character by looking up in a $G \times p$ character-embedding matrix, where p is the dimensions of the character-embedding. The sequence of character-embedding vectors, e_char_t , of the word is then fed to a bidirectional LSTM [12]. The final character representation, e'_char_t , is obtained by concatenating the forward and backward final states.

The final word embedding vector e_t for each word is thus,

$$e_t = [e_word_t, e'_char_t] \quad (3)$$

3.3 Encoder

We use a single layer Bi-Directional RNN as the encoder, with LSTM as the basic recurrent unit due to its ability to incorporate long-term dependencies [12]. As the purpose of the encoder is to capture the shared representations of the multiple tasks, these representations should include contextual information in both the forward and backward directions, where the outputs of each task can depend on these previous and future elements in the sequence. So we pass the input sequence through a Bi-LSTM to serve this purpose. The Bi-LSTM achieves this by passing the input sequence in its original order through a forward LSTM, which encodes a hidden state $\vec{h}_t$ for each time step t . Also, a reversed copy of the input sequence is passed through a backward LSTM, that encodes a hidden state $\overleftarrow{h}_t$. The forward and backward hidden states are then concatenated to represent the final hidden state of the encoder at each time step.

$$h_t = [\vec{h}_t, \overleftarrow{h}_t] \quad (4)$$

$$\vec{h}_t = LSTM[x_t, \vec{h}_{t-1}] \quad (5)$$

$$\overleftarrow{h}_t = LSTM[x_t, \overleftarrow{h}_{t-1}] \quad (6)$$

3.4 Decoder

On the decoder side, we allocate a single layer uni-directional LSTM with attention mechanism for each task in order to produce output specifically for that task. The *Coverage* mechanism is integrated in the attention mechanism to give the model a sense of how much attention it has already assigned so far.

3.4.1 Attention Mechanism. Conventional Encoder-Decoder for sequence-to-sequence tasks has shown to have limitations as it tries to decode a fixed-length encoded vector when generating output [2]. The fact that the last encoder hidden state is required to hold the summary of all the timesteps, in practice, doesn't hold, especially, for longer sentences [6].

The use of attention has come a long way in tasks ranging from image captioning [36] to machine translation [6]. By peeking over all the encoder states and weighing them according to their relevance to the current output to generate, it produces the *attention distribution*. This attention distribution gives a signal where to attend to more in the input sequence.

We use additive attention from [2], where *attention distribution* a_t is calculated as:

$$g_{ti} = v_a^T \tanh(W_a h_i + W_b s_{t-1} + b_{attn}) \quad (7)$$

$$a_t = \text{softmax}(g_t) \quad (8)$$

h_i is an encoder hidden state and s_{t-1} is the previous decoder hidden state at decoding timestep t . v_a , W_a , W_b , and b_{attn} are learnable attention parameters. Finding an attention score using the attention function $g_{ti}(h_i, s_{t-1})$ between each encoder state h_i and the decoder state preceding the current output state, tells exactly where to pay attention to in the input sequence when decoding.

The encoder hidden states are then combined with the attention distribution to give the *context vector* defined as the attention weighted sum of the encoder hidden states,

$$con_t = \sum_i h_i a_{ti} \quad (9)$$

This context vector, along with the previous decoder hidden state s_{t-1} , outputs the label generated at previous timestep, y_{t-1} . The aligned encoder hidden state h_t is used to compute the decoder state s_t at time step t for ADR and Indication labeling tasks.

$$s_t = f(s_{t-1}, h_t, y_{t-1}, con_t) \quad (10)$$

Whereas for the ADR Classification task, the decoder state at time step t is a function of s_{t-1} and con_t .

$$s_t = f(s_{t-1}, con_t) \quad (11)$$

For all the tasks, initial decoder hidden state s_0 is set to the final encoder hidden state.

3.4.2 Coverage Mechanism. Due to the colloquial nature of conversations on social network, ADRs are expressed in everyday language and also take phrasal form. Consider the following example tweet, *Just took Seroquel. Now I'm freaking out that I will end up sleeping 15 hrs and miss my 12pm appt tomorrow*, where *sleeping 15 hrs* is an ADR taking a phrase form. The phrasal ADRs can also be expressed as a list of ADRs, illustrated by the tweet *@user I hated Effexor. It makes you hungry, dizzy, and lethargic. That culminated in a large weight gain for me*, where *hungry*, *dizzy*, and *lethargic* are a list of ADRs occurring as a phrase for the drug Effexor. As we found from preliminary experiments that with just attention it cannot detect all the ADR words in these phrases, we introduce *coverage* into our decoder to solve this problem. For every decoder timestep we keep track of a coverage vector c_t , implemented as in [33] for summarization. The coverage vector sums up the attention distribution from all the previous decoder timesteps. However, unlike in [33] we impose a window for the previous decoder timesteps and set it to 3. We do this as most ADRs are comprised of few words and considering attention from all the previous timesteps can include attention from non-ADR words, which can jeopardize the attention distribution of ADR phrase located at the end of a post. With the coverage vector, in case if a word in a phrasal ADR does not get identified, the attentive words based on the attention distribution for the neighboring ADR words can help locate it. In other words, it helps secure the ADR label for the other words in a phrasal ADR. Coverage also contributes in avoiding mislabeling an Indication as an ADR and vice versa as it implicitly keeps track of the location of a word. That is, unlike conventional additive attention - which only

uses encoder states h_t to find the attentive words - by summing up attention values assigned so far, it provides information about how far the model has attended from the beginning of the sentence. This way it learns the boundaries of attention for the ADR and Indication words. It is defined as,

$$c_t = \sum_{s=t-1}^{s=t-3} a_s \quad (12)$$

With only attention mechanism, if for the word *sleeping* in the first example attention was focused on the words *Seroquel*, *freaking* and *out* to identify it as an ADR, it does not guarantee that it will attend to similar words for *15* and recognize it as a part of the ADR phrase since *15* can frequently appear with words irrelevant to ADR mention. By augmenting attention mechanism with *coverage*, we hypothesize that by also focusing on the words that were attended to in the previous time steps, *15* can be properly tagged as an ADR. That is, when decoding *15* it will pay attention to the words *Seroquel*, *freaking* and *out*, that were helpful in tagging *sleeping* correctly. Hence this hybrid attention with coverage module can enable capturing the words in an ADR phrase, which otherwise just with attention can get missed.

We update the attention distribution to the following as coverage is now part of the attention mechanism,

$$g_{ti} = v_a^T \tanh(W_a h_i + W_b s_{t-1} + W_c c_{ti} + b_{attn}) \quad (13)$$

3.5 Output

Our task-specific output is calculated as the following.

ADR Classification: The context vector can be viewed as a high-level representation of the post and is used as features for the classification task.

$$y = \text{softmax}(W_{class} con + b_{class}), \quad (14)$$

where W_{class} and b_{class} are learnable parameters.

ADR Labeling and Indication Labeling: The context vector at each decoder timestep, con_i , can be viewed as the final representation for the word, which is used with the decoder state, s_i , to predict the output at that timestep.

$$y_i = \text{softmax}(W' [s_i, con_i] + b'), \quad (15)$$

where W' and b' are learned during training independently for the ADR labeling and Indication labeling tasks.

4 TRAINING THE MODEL

We train our model jointly over all the datasets. Cross-entropy is used as the loss function during training, which is defined for each task as:

ADR Classification:

$$L(\theta_{class}) = -\frac{1}{n} \sum_x (Y_x \log Y'_x + (1 - Y_x) \log(1 - Y'_x)), \quad (16)$$

where n is the training data size, x means over all inputs, Y is the true label and Y' the predicted label probability. $\theta_{class} = (\theta_{src}, \theta_1)$, where θ_{src} is a collection of shared parameters among all the tasks for the encoder and θ_1 are the parameters for ADR Classification decoder.

ADR Labeling:

$$L(\theta_{ADRLab}) = \frac{1}{n} \sum_x \log P(Y_x | X_x; \theta_{ADRLab}), \quad (17)$$

where $\theta_{ADRLab} = (\theta_{src}, \theta_2)$ and θ_2 is the parameter for ADR Labeling decoder.

Indication Labeling:

$$L(\theta_{INDlab}) = \frac{1}{n} \sum_x \log P(Y_x | X_x; \theta_{INDlab}), \quad (18)$$

where $\theta_{INDlab} = (\theta_{src}, \theta_3)$ and θ_3 is the parameter for Indication Labeling decoder.

During training, we use the weighted loss which consists of a total number of T tasks:

$$\text{total loss} = \sum_{t=1}^T \omega_t L_t. \quad (19)$$

Total loss is a linear combination of loss for all tasks, where ω_t is the weight for each task t respectively.

5 EXPERIMENTAL SETTINGS

5.1 Datasets

ADR Classification: We use the Twitter dataset from PSB 2016 Social Media Shared Task for ADR Classification [32]. This dataset was created from tweets collected using generic and brand names of the drugs, along with their phonetic misspellings. It contains binary annotations of '0' and '1' referring to 'ADR' and 'notADR' respectively. Although the original dataset contains a total of 10,822 tweets, we could download only 7044 tweets from the given IDs. We consider only the 1081 tweets out of these that overlap with the tweets in ADR Labeling dataset as belonging to 'ADR' class. Another 1081 tweets are randomly sampled from the 'notADR' class. This is done as all the tasks share the same encoder and we exploit the shared input representations. We manually labeled the supplemental dataset discussed below and add it to the existing dataset. We divided the tweets randomly into training, test and validation datasets with splits of 70%-15%-15%.

ADR Labeling: For ADR Labeling, we use the Twitter dataset from PSB 2016 Social Media Shared Task for ADR Extraction [32]. It contains around 2000 tweets annotated with tweet ID, start offset, end offset, semantic type (ADR/Indication), UMLS ID, annotated text span and the related drug. However, at the time of this study 1081 of the annotated tweets were available for download. We supplemented this dataset with a small dataset [7] of 203 tweets collected between May 2015 to December 2015 from Twitter. We use the split 70%-15%-15% for the training, test and validation datasets. We customized this dataset to only include the ADR annotations.

Indication Labeling: The corpus and splits used for indication labeling are the same as that for ADR Labeling. We customized this dataset to only include the Indication annotations.

5.2 Training Details

Our implementation is based on the open source deep learning package Tensorflow [1]. Glove [28] toolkit was used to pre-train the word embeddings, which were then used to initialize the embeddings in the model. We specifically used the model trained on Twitter in Glove to be able to generate the word-embeddings of

the words unique to our Twitter dataset. We set the number of units in the LSTM cell to 128 and the dimensionality of word and char embeddings are set to 200 and 128 respectively. All the forget biases in the LSTMs are set to 1. In every epoch, we perform mini-batch training of each parallel task corpus with a batch size of 16. Regularization is done on the non-recurrent connections with a dropout rate of 0.5. We used Adam Optimization [15] method with a learning rate of 0.1 during training. All weights and biases in the attention and coverage component were initialized with Xavier [11] initialization. The development set was used to tune the value of the hyper-parameter weight, ω , for the loss of each task. The setting 1, 1 and 0.1 for the losses of ADR labeling, Indication labeling, and ADR classification are used for the final experiments as it gave best results on the development set.

6 RESULTS AND DISCUSSION

6.1 Evaluation

We use precision, recall and F-1 score as the evaluation measures. For a particular class i , precision and recall can be defined with the respective equations $p_i = \frac{TP_i}{TP_i + FN_i}$ and $r_i = \frac{TP_i}{TP_i + FP_i}$, where TP_i is the number of true positives, FN_i is the number of false negatives and FP_i is the number of false positives. F-1 score is calculated from the precision and recall as $F-1 = \frac{2p_i r_i}{p_i + r_i}$.

For two tagging tasks, namely, ADR and Indication Labeling, we did approximate matching [7, 35] on the predicted labels for the ADR phrasal words against their actual labels. Approximate matching works by checking if one or more of the ADR spans in an ADR phrase could be identified correctly with the ‘ADR’ tags. For example, for the following tweet *I was on Cymbalta for 5 days. Cold turkey had sweats, migraine, tremors while on & 3 days after.* with the actual ADR span *sweats, migraine, tremors*, predicting the tag as ‘ADR’ for any of the three spans or their combinations would be considered correct. The approximate match precision and recall are calculated as:

$$p_i = \frac{\text{\# of ADR spans correctly tagged}}{\text{\# of predicted ADR spans}} \quad (20)$$

$$r_i = \frac{\text{\# of ADR spans correctly tagged}}{\text{\# of actual ADR spans}} \quad (21)$$

6.2 Baselines

As we could not find any previous work that performs multi-task learning on pharmacovigilance tasks, we compare it against two baseline methods and state-of-the-art approaches for the independent tasks to demonstrate the effectiveness of our proposed model.

- **BLSTM-Random, BLSTM-Pretrained-learnable, BLSTM-Pretrained-fixed:** The architecture of this model is known as Bidirectional Long Short-Term Memory (BLSTM) RNN [7]. It combines a forward RNN and a backward RNN and uses only word embeddings as the features. In BLSTM-Random the word-embeddings are randomly initialized and treated as learnable parameters. BLSTM-Pretrained-learnable and BLSTM-Pretrained-fixed use pre-trained word-embeddings trained on a large non-domain specific Twitter dataset. The

only difference between them is that BLSTM-Pretrained-learnable treats the words-embedding values as learnable parameters, while BLSTM-Pretrained-fixed as fixed constants.

- **CRNN:** State-of-the-art model for doing ADR Classification task. CRNN [13] is a convolutional neural network concatenated with a recurrent neural network. They used GRU as the basic RNN unit and RLU for the convolutional layer.
- **CNNA:** State-of-the-art model for doing ADR Classification task. CNNA [13] is a convolutional neural network with attention mechanism incorporated.
- **MT-NoAtten:** Our multi-task framework for pharmacovigilance tasks without any attention. We use a non-attention RNN as the decoder in this case.
- **MT-Atten:** Our multi-task framework for pharmacovigilance tasks with just attention mechanism. Coverage is turned off during training.
- **MT-Atten-Cov:** This is our proposed multi-task framework for pharmacovigilance tasks with combined attention and coverage mechanism.

6.3 Overall Performance

To show the validity of our model, we report results on three experiments. The results obtained from applying our model to the test sets are presented in Tables 1, 2 and 3.

In the first experiment, we train our multi-task model jointly on the three parallel datasets corresponding to each of the three tasks. Results for each task from this experiment are compared against two baselines, which are depicted in Table 1. As the purpose of coverage is to be able to have a greater coverage for the ADR/ Indication words in an ADR/ Indication phrase, we turn it off in the classification decoder and use it in the two tagging tasks. So, results for classification task are not reported with MT-Atten-Cov model. For the remaining two tasks, we can observe that our method outperforms both the baselines in terms of the precision, recall, and F-score. Although approximate matching would consider the identification of any ADR/ Indication span as a true positive and we would expect comparable results for MT-Atten-Cov and MT-Atten models, the fact that MT-Atten-Cov has superior results empirically confirms that incorporating coverage with attention helps in capturing ‘phrasal’ and single ADR words which are not attended to with just the attention mechanism. We achieve a 1.50% and 0.90% F-1 score improvement for ADR labeling and Indication labeling respectively over the MT-Atten model. The experimental values for Indication labeling task happen to be low across all the models due to the sparsity of tweets containing indication words. Nevertheless, we gain improvement with our model against both the baselines. We gain an F-1 score improvement of 21.32%, 13.59% and 23.43% over the MT-NoAtten model for classification, ADR detection, and Indication labeling respectively. By examining the precision and recall results, we can say that the better performance of our model can be attributed to the improvement of both.

In the second and third experiments, we train our model separately on ADR classification and ADR detection tasks respectively and provide the comparison against results from the model trained in experiment one. As Indication detection was not performed as an independent task in any previous works, we do not provide

Table 1: Test set results for the three tasks with the proposed model and the two baselines.

Metrics	MT-Atten-Cov			MT-Atten Baseline			MT-NoAtten Baseline		
	P (%)	R (%)	F-1 (%)	P (%)	R (%)	F-1 (%)	P (%)	R (%)	F-1 (%)
ADR Classification	N/A	N/A	N/A	72.88	70.54	70.69	69.63	60.60	55.62
ADR Labeling	72.31	87.5	79.24	70.88	86.81	78.04	60.50	78.88	68.47
Indication Labeling	47.50	50.2	48.82	46.87	50.00	48.38	34.22	41.20	37.38

Table 2: Comparison of ADR Classification Task test results to previous approaches. Single-Atten-Cov Task refers to the independent task model trained on only the ADR Classification dataset.

	P (%)	R (%)	F-1 (%)
CRNN	49.00	55.00	51.00
CNNA	40.00	66.00	49.00
Single-Atten-Cov	70.21	70.05	70.13
MT-Atten	72.88	70.54	70.69

Table 3: Comparison of ADR Labeling Task test results to previous approaches. Single-Atten-Cov Task refers to the independent task model trained on only the ADR Labeling dataset.

	P (%)	R (%)	F-1 (%)
BLSTM-Random	64.57	63.32	62.72
BLSTM-Pretrained-learnable	60.47	80.70	68.58
BLSTM-Pretrained-fixed	70.43	82.86	75.49
Single-Atten-Cov	71.50	86.22	78.17
MT-Atten-Cov	72.31	87.50	79.24

a separate table with comparisons to previous approaches for it. We can see from Table 2 that our multi-task model improves the performance in terms of the F-1 score with 0.79 % for classification compared to a single task model. While for the ADR detection task results in Table 3, it makes 1.35% improvement. These empirical findings cast light on that shared input representations and interactions between tasks result in mutual benefit for all tasks. Comparing all the independent classification models (Single-Atten-Cov Task, CRNN, and CNNA) against each other, we can further see the advantage of using attention for the classification task, where Single-Atten-Cov has an average improvement of 28.71% over CRNN and CNNA. Although CNNA incorporates attention in its model, we assume using an RNN encoder-decoder with attention is more helpful. Similarly for ADR detection, from Table 3 we can see that our single task ADR detection model (Single-Atten-Cov Task) outperforms the best model among BLSTM (BLSTM-Pretrained-fixed) by 3.43%. Classification performance for our multi-task model attained higher F-1 scores over the CRNN and CNNA models by 27.85% and 30.66% respectively, while for ADR detection, multi-task model improves by 4.73% over the best performing BLSTM models.

6.4 Case Study

To get a deeper insight into how augmenting coverage to attention mechanism benefits our model, we sample several tweets from the test dataset. The tags predicted by our model for these tweets are compared against those from the baseline MT-Atten model. All usernames are anonymized for privacy concerns. In order to validate the results that our model produced, we further visualize the attention heatmaps for some of the tweets depicted in Figure 3.

The following two tweets illustrate the ability of our model to correctly label the single ADR word as ‘ADR’, while the MT-Atten model makes wrong prediction. This justifies the higher precision and recall gained by our model over MT-Atten.

- **Tweet 1:** @user1 bloody zombie I also take Venlafaxine thats
for cronic depression bipolar is a difficult illness to deal with
- True Tags:** [‘O’, ‘O’, ‘ADR’, ‘O’, ‘O’, ‘O’, ‘O’, ‘O’, ‘O’, ‘O’,
‘O’, ‘O’, ‘O’, ‘O’, ‘O’, ‘O’, ‘O’, ‘O’]
- MT-Atten:** [‘O’, ‘O’, ‘O’, ‘O’, ‘O’, ‘O’, ‘O’, ‘O’, ‘O’, ‘O’, ‘ADR’,
‘O’, ‘O’, ‘O’, ‘ADR’, ‘O’, ‘O’, ‘O’, ‘O’]
- MT-Atten-Cov:** [‘O’, ‘O’, ‘ADR’, ‘O’, ‘O’, ‘O’, ‘O’, ‘O’, ‘O’, ‘O’,
‘O’, ‘O’, ‘O’, ‘O’, ‘O’, ‘O’, ‘O’, ‘O’]
- ADR:** zombie

The degree of *attention* received by each source word when decoding a target word for Tweet 1 is visualized in Figure 3, where darker shades denote higher score. When it is predicting the tag for target token ‘zombie’, more attention is given to the words ‘zombie’, ‘Venlafaxine’, ‘difficult’ and ‘illness’. Due to coverage, attentive words ‘difficult’ and ‘illness’ from previous timesteps were also attended. These words facilitated in predicting the ‘ADR’ tag for this word. Furthermore, *indication* words ‘bipolar’ and ‘depression’ receiving low attention weights are prevented from mislabeling.

- **Tweet 2:** *Cymbalta*, my mood has worsened
True Tags: ['O', 'O', 'ADR', 'O', 'O']
MT-Atten: ['O', 'O', 'O', 'O', 'ADR']
MT-Atten-Cov: ['O', 'O', 'ADR', 'O', 'ADR']
ADR: mood

The following tweets contain ADRs captured by our model but missed by MT-Atten.

- **Tweet 3:** *I'm sorry you have ostrich halitosis @user2 but that IS a side effect of once monthly boniva.*
- True Tags:** ['O', 'O', 'O', 'O', 'O', 'O', 'ADR', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O']
- MT-Atten:** ['O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O']
- MT-Atten-Cov:** ['O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'ADR', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O']
- ADR:** halitosis

Figure 3: Attention Heatmaps for Tweets 1, 4 and 5.

- **Tweet 4:** Crying randomly at nothing and everything. Sigh. Thank you #effexor #withdrawals Can you just exit my system already.

True Tags: [ADR, 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'ADR', 'O', 'O', 'O', 'O', 'O', 'O', 'O']

MT-Atten: ['O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O']

MT-Atten-Cov: ['O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O']

ADR: Crying ADR: withdrawals

From the attention heatmap of this tweet shown in Figure 3, we can see that the first ADR word 'crying' has overall low attention for all the words. Also, as it happens to be the first word of the sentence, no coverage was passed. This is the reason why it is not detected as an ADR by both the baseline and our model. Whereas, for the second ADR 'withdrawals' source words 'crying', 'randomly', 'at' and 'nothing' were attended to more than others. Moreover, 'Effexor' and 'withdrawals' from the attention of previous word 'Effexor' were also instrumental in the right tag assignment of 'withdrawals'.

With the following tweet we can see how *coverage* causes attention from previous ADR word 'heart' to identify 'was' as part of the ADR phrase. We can see the baseline model MT-Atten mislabels it.

- **Tweet 5:** @user3 I took a vyvanse and drank a 12oz redbull this morning. My heart was hurting a lil bit haha

True Tags: ['O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'ADR', 'ADR', 'ADR', 'O', 'O', 'O', 'O']

MT-Atten: ['O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'O', 'ADR', 'O', 'ADR', 'O', 'O', 'O', 'O']

MT-Atten-Cov: ['O', 'ADR', 'ADR', 'ADR', 'O', 'O', 'O', 'O']

ADR: heart was hurting

7 RELATED WORK

Pharmacovigilance in Social Media: Pharmacovigilance has become a very active area of research [24]. Initial work on pharmacovigilance in social media performed detection and extraction tasks on health forum (e.g., DailyStrength, MedHelp) data [17, 21, 25]. These works also focused on investigating posts mentioning adverse reactions associated with a fewer number of drugs. Over the years, a wide range of drugs have been used and other platforms such as Twitter have emerged as a valuable source for ADR monitoring [9, 10, 23].

Two of the most prevalent tasks for pharmacovigilance studies are the detection of posts containing ADR mention and extraction of the ADR mentions from posts [31]. Due to the scarcity of annotated resources for this study, most works have been performed using small annotated datasets [17, 21]. While some work used large

unannotated dataset employing an unsupervised approach [3, 5]. For ADR extraction, ADR lexicons and knowledge bases have been the most widely used resource [31]. However, medical terms are rarely used in social media posts and raise mapping issues. Applying deep neural networks for pharmacovigilance has found its way in some recent works [7, 13]. These approaches benefit from not having to explicitly specify the features, but rather learn them in the training process.

As mention of ‘Indication’ co-occurs with ADRs in a post, no previous work has performed ‘Indication’ detection/extraction as a separate task. Although some works have used it as features in a machine learning approach [22], the sparsity of ‘indications’ has been an issue. With our multi-task approach that includes ‘Indication’ extraction as a separate task, we show that interaction between related tasks can improve learning for all.

Multi-task Learning: Use of multi-task learning models has become ubiquitous for many machine learning applications in areas ranging from natural language processing, speech recognition to computer vision [30]. Multi-task learning with encoder-decoder come in one of the three flavors: one-to-many, many-to-one and many-to-many approaches [19]. In One-to-many models, tasks share a common encoder and have a task-specific decoder. A one-to-many approach was used in [8] for translation task from a source language to multiple target languages. The shared encoder captured the syntactic and semantic similarity existing across different languages. On the other hand, a many-to-one approach is suitable for tasks where a decoder can be easily shared, such as multi-source translation. Lastly, a many-to-many approach allows multiple encoders and decoders.

Our one-to-many approach is similar in spirit to [18] used for joint slot filling and intent detection tasks. However, our work has the following differences: 1) learns more tasks jointly 2) use of coverage with attention mechanism. 3) use of pre-trained word embeddings.

8 CONCLUSION

Performing pharmacovigilance on Twitter can augment the existing ADR surveillance system which suffers from various clinical limitations. In this work, we provide an end-to-end solution for three ADR detection and extraction tasks. These problems have conventionally been approached separately and didn’t leverage the interactions between the tasks. Exploiting the similarity that these tasks have, we proposed a multi-task encoder-decoder framework. All the tasks share one encoder to model the interactions and semantic/syntactic similarity between them. While each has its own decoder to produce output specific to that task. Our empirical findings validate how learning all tasks jointly improves precision and recall over state-of-the-art approaches. Additionally, the proposed solution is a hybrid attention model with coverage to deal with ADRs occurring as phrases. Through results and case studies we show that not only is this hybrid model able to achieve higher phrasal ADR word coverage compared to the baselines but is also able to identify single ADRs correctly.

9 ACKNOWLEDGMENTS

This work is supported in part by NSF through grants IIS-1526499, CNS-1626432 and NSFC 61672313.

REFERENCES

- [1] Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. TensorFlow: A System for Large-Scale Machine Learning. In *OSDI*, Vol. 16. 265–283.
- [2] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473* (2014).
- [3] Adrian Benton, Lyle Ungar, Shawndra Hill, Sean Hennessy, Jun Mao, Annie Chung, Charles E Leonard, and John H Holmes. 2011. Identifying potential adverse effects using the web: A new approach to medical hypothesis generation. *Journal of biomedical informatics* 44, 6 (2011), 989–996.
- [4] Jiang Bian, Umit Topaloglu, and Fan Yu. 2012. Towards large-scale twitter mining for drug-related adverse events. In *Proceedings of the 2012 international workshop on Smart health and wellbeing*. ACM, 25–32.
- [5] Brant W Chee, Richard Berlin, and Bruce Schatz. 2011. Predicting adverse drug events from personal health messages. In *AMIA Annual Symposium Proceedings*, Vol. 2011. American Medical Informatics Association, 217.
- [6] Kyunghyun Cho, Bart Van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259* (2014).
- [7] Anne Cocos, Alexander G Fiks, and Aaron J Masino. 2017. Deep learning for pharmacovigilance: recurrent neural network architectures for labeling adverse drug reactions in Twitter posts. *Journal of the American Medical Informatics Association* (2017), ocw180.
- [8] Daxiang Dong, Hua Wu, Wei He, Dianhai Yu, and Haifeng Wang. 2015. Multi-Task Learning for Multiple Language Translation. In *ACL (1)*. 1723–1732.
- [9] Clark C Freifeld, John S Brownstein, Christopher M Menone, Wenjie Bao, Ross Filice, Taha Kass-Hout, and Nabarun Dasgupta. 2014. Digital drug safety surveillance: monitoring pharmaceutical products in twitter. *Drug safety* 37, 5 (2014), 343–350.
- [10] Rachel Ginn, Pranoti Pimpalkhute, Azadeh Nikfarjam, Apurv Patki, Karen O’Connor, Abeed Sarker, Karen Smith, and Graciela Gonzalez. 2014. Mining Twitter for adverse drug reaction mentions: a corpus and classification benchmark. In *Proceedings of the fourth workshop on building and evaluating resources for health and biomedical text processing*.
- [11] Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*. 249–256.
- [12] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [13] Trung Huynh, Yulan He, Alistair Willis, and Stefan Rüger. 2016. Adverse drug reaction classification with deep neural networks. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*. 877–887.
- [14] Keyuan Jiang and Yujing Zheng. 2013. Mining Twitter data for potential drug effects. In *International Conference on Advanced Data Mining and Applications*. Springer, 434–443.
- [15] Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [16] Vasileios Lampsos, Elad Yom-Tov, Richard Pebody, and Ingemar J Cox. 2015. Assessing the impact of a health intervention via user-generated Internet content. *Data Mining and Knowledge Discovery* 29, 5 (2015), 1434–1457.
- [17] 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 *Proceedings of the 2010 workshop on biomedical natural language processing*. Association for Computational Linguistics, 117–125.
- [18] Bing Liu and Ian Lane. 2016. Attention-based recurrent neural network models for joint intent detection and slot filling. *arXiv preprint arXiv:1609.01454* (2016).
- [19] Minh-Thang Luong, Quoc V Le, Ilya Sutskever, Oriol Vinyals, and Lukasz Kaiser. 2015. Multi-task sequence to sequence learning. *arXiv preprint arXiv:1511.06114* (2015).
- [20] Ramesh Nallapati, Feifei Zhai, and Bowen Zhou. 2017. SummaRuNNer: A recurrent neural network based sequence model for extractive summarization of documents. *hiP (yi= 1/ hi, si, d) 1* (2017), 1.
- [21] Azadeh Nikfarjam and Graciela H Gonzalez. 2011. Pattern mining for extraction of mentions of adverse drug reactions from user comments. In *AMIA Annual Symposium Proceedings*, Vol. 2011. American Medical Informatics Association, 1019.
- [22] Azadeh Nikfarjam, Abeed Sarker, Karen O’Connor, Rachel Ginn, and Graciela Gonzalez. 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 (2015), 671–681.
- [23] Karen O’Connor, Pranoti Pimpalkhute, Azadeh Nikfarjam, Rachel Ginn, Karen L Smith, and Graciela Gonzalez. 2014. Pharmacovigilance on twitter? mining tweets

- for adverse drug reactions. In *AMIA annual symposium proceedings*, Vol. 2014. American Medical Informatics Association, 924.
- [24] World Health Organization et al. 2002. The importance of pharmacovigilance. (2002).
- [25] Apurv Patki, Abeer Sarker, Pranoti Pimpalkhute, Azadeh Nikfarjam, Rachel Ginn, Karen O'Connor, Karen Smith, and Graciela Gonzalez. 2014. Mining adverse drug reaction signals from social media: going beyond extraction. *Proceedings of BioLinkSig 2014* (2014), 1–8.
- [26] Michael J Paul and Mark Dredze. 2011. You are what you Tweet: Analyzing Twitter for public health. *ICWSM 20* (2011), 265–272.
- [27] Mor Peleg, Tiffany I Leung, Manisha Desai, and Michel Dumontier. 2017. Is Crowdsourcing Patient-Reported Outcomes the Future of Evidence-Based Medicine? A Case Study of Back Pain. In *Conference on Artificial Intelligence in Medicine in Europe*. Springer, 245–255.
- [28] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. 1532–1543.
- [29] Marek Rei, Gamal KO Crichton, and Sampo Pyysalo. 2016. Attending to Characters in Neural Sequence Labeling Models. *arXiv preprint arXiv:1611.04361* (2016).
- [30] Sebastian Ruder. 2017. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098* (2017).
- [31] Abeer Sarker, Rachel Ginn, Azadeh Nikfarjam, Karen O'Connor, Karen Smith, Swetha Jayaraman, Tejaswi Upadhaya, and Graciela Gonzalez. 2015. Utilizing social media data for pharmacovigilance: A review. *Journal of biomedical informatics* 54 (2015), 202–212.
- [32] Abeer Sarker, Azadeh Nikfarjam, and Graciela Gonzalez. 2016. Social media mining shared task workshop. In *Biocomputing 2016: Proceedings of the Pacific Symposium*. 581–592.
- [33] Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get To The Point: Summarization with Pointer-Generator Networks. *arXiv preprint arXiv:1704.04368* (2017).
- [34] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*. 3104–3112.
- [35] Richard Tzong-Han Tsai, Shih-Hung Wu, Wen-Chi Chou, Yu-Chun Lin, Ding He, Jieh Hsiang, Ting-Yi Sung, and Wen-Lian Hsu. 2006. Various criteria in the evaluation of biomedical named entity recognition. *BMC bioinformatics* 7, 1 (2006), 92.
- [36] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2015. Show and tell: A neural image caption generator. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3156–3164.
- [37] Stephen Wu, Sijia Liu, Yanshan Wang, Tamara Timmons, Harsha Uppili, Steven Bedrick, William Hersh, and Hongfang Liu. 2017. Intrainstitutional EHR collections for patient-level information retrieval. *Journal of the Association for Information Science and Technology* (2017).
- [38] Zhengzheng Xing, Jian Pei, and Eamonn Keogh. 2010. A brief survey on sequence classification. *ACM Sigkdd Explorations Newsletter* 12, 1 (2010), 40–48.
- [39] Christopher C Yang, Haodong Yang, Ling Jiang, and Mi Zhang. 2012. Social media mining for drug safety signal detection. In *Proceedings of the 2012 international workshop on Smart health and wellbeing*. ACM, 33–40.
- [40] Ming Yang, Xiaodi Wang, and Melody Y Kiang. 2013. Identification of Consumer Adverse Drug Reaction Messages on Social Media.. In *PACIS*. 193.
- [41] Elad Yom-Tov. 2017. Predicting drug recalls from Internet search engine queries. *IEEE Journal of Translational Engineering in Health and Medicine* 5 (2017), 1–6.
- [42] Bin Zou, Vasileios Lampos, Russell Gorton, and Ingemar J Cox. 2016. On infectious intestinal disease surveillance using social media content. In *Proceedings of the 6th International Conference on Digital Health Conference*. ACM, 157–161.