

Want to Identify, Extract and Normalize Adverse Drug Reactions in Tweets? Use RoBERTa

Katikapalli Subramanyam Kalyan
Department of Computer Applications
NIT Trichy, India
kalyan.ks@yahoo.com

Sivanesan Sangeetha
Department of Computer Applications
NIT Trichy, India
sangeetha@nitt.edu

Abstract

This paper presents our approach for task 2 and task 3 of Social Media Mining for Health (SMM4H) 2020 shared tasks. In task 2, we have to differentiate adverse drug reaction (ADR) tweets from nonADR tweets and is treated as binary classification. Task 3 involves extracting ADR mentions and then mapping them to MedDRA codes. Extracting ADR mentions is treated as sequence labeling and normalizing ADR mentions is treated as multi-class classification. Our system is based on pre-trained language model RoBERTa and it achieves a) F1-score of 58% in task 2 which is 12% more than the average score b) relaxed F1-score of 70.1% in ADR extraction of task 3 which is 13.7% more than the average score and relaxed F1-score of 35% in ADR extraction + normalization of task 3 which is 5.8% more than the average score. Overall, our models achieve promising results in both the tasks with significant improvements over average scores.

1 Introduction

Social media platforms in particular, twitter are used extensively by common public to share their experiences which also includes health-related information like adverse drug reaction (ADR) they experience while consuming drugs. Adverse Drug Reaction refers to unwanted harmful effect following the use of one or more drugs. The abundant health-related social media data can be utilized to enhance the quality of services in health-related applications (Kalyan and Sangeetha, 2020b). Our team participated in task 2 and task 3 of SMM4H 2020 shared task (Klein et al., 2020). Task 2 aims at identifying whether a tweet contains ADR mention or not. Task 3 aims at extracting ADR mentions and then normalizing them to MedDRA concepts.

2 Task 2 – Identification of ADR tweets

2.1 Problem Definition and Dataset

Task 2 aims at identifying whether a tweet contains ADR mention or not. An example of an ADR tweet is: *‘thank god for vyvanse #addicted’*. Here *‘addicted’* is the adverse drug reaction that happened because of consumption of the drug *‘vyvanse’*. An example of a nonADR tweet is *‘never take paxil #js’*. In this task, we learn a classification model which outputs the label 1 or 0 for a given tweet depending on whether it contains ADR mentions or not. In this dataset, the training set consists of 20544 tweets (18641 nonADR and 1903 ADR tweets), validation set consists of 5134 tweets (4660 nonADR and 474 ADR tweets) and the test set consists of 4759 tweets.

2.2 Methodology

Pre-processing

We apply the following pre-processing steps

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

- Lowercase the text and remove consecutively repeating characters in the words (e.g., feeeeel → feel).
- Remove urls, @user mentions, retweet tag (rt), non-ASCII and punctuation characters.
- Expand English contractions (e.g., can't → cannot) and replace interjections with their meanings (e.g., ouch, oww → pain).
- Replace character smiley (e.g., :) → happy) and emoji (e.g., 😊 → grinning face) with their text descriptions.

Model Description

In recent times, the evolution of pretrained language models like BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019) changed the scenario in natural language processing from training downstream models from scratch to just fine-tuning pre-trained models. These models learn universal language representations from large training corpus and they can be used in downstream tasks by adding one or two layers which are specific to the task (Qiu et al., 2020). Our model is based on RoBERTa. We add task-specific sigmoid layer on the top of RoBERTa and then fine-tune the entire model using training dataset. We consider the final hidden state vector $t_{<s>} \in \mathbb{R}^h$ of the special token $<s>$ as tweet representation. Here h represents hidden state vector size in RoBERTa-base and it is equal to 768. The vector $t_{<s>}$ is passed through sigmoid layer to get the prediction \hat{y} . Overall, the label \hat{y} is computed as :

$$t_{<s>} = RoBERTa(tweet) \quad (1)$$

$$\hat{y} = Sigmoid(t_{<s>}W^T + b) \quad (2)$$

Here $W \in \mathbb{R}^{1 \times h}$ and $b \in \mathbb{R}$ are learnable parameters of sigmoid layer.

2.3 Experiments and Results

As ADR tweets are less in number compared to nonADR tweets in the training set, we augment training set with ADR tweets from SMM4H 2017 (Sarker et al., 2018) and SMM4H 2019 (Weissenbacher et al., 2019) ADR tweets classification datasets. Further, we include only randomly chosen 90% of nonADR tweets in the training set. By conducting random search over the range of hyperparameters values, we arrive at the following optimal set of hyperparameter values: batch size = 128, learning rate = 3e-5, dropout = 0.2 (applied on $t_{<s>}$ vector to reduce overfitting) and epochs = 10. We implement our model in PyTorch framework using transformers library from huggingface (Wolf et al., 2019).

| Model | Precision | Recall | F1 |
|----------------|-----------|--------|--------------|
| RoBERTa (ours) | 52.00 | 65.00 | 58.00 |
| Average scores | 42.00 | 59.00 | 46.00 |

Table 1: Task 2 - ADR Tweet classification results on test data

We report performance of our model and average scores in task 2 – ADR tweets classification in Table 1. Our model achieves an F1-score of 58% and it is 12% more than the average score.

3 Task 3 – Extract and Normalize ADR Mentions

3.1 Problem Definition and Dataset

This task involves ADR extraction followed by normalization. In the first part (ADR Extraction), to extract ADR mentions the model has to identify ADR tweets and then extract ADR mentions by identifying the spans in tweets. A tweet can have more than one ADR mention also and an ADR mention can be a sequence of words also. Example of a tweet with ADR mentions

@coolpharmgreg i don't care if they are toxic haha putting the cipro drops in is essentially equivalent to torture #oww

Here ‘*oww*’, ‘*toxic*’ and ‘*equivalent to torture*’ are the adverse drug reactions due to the consumption of the drug ‘*cipro*’. In the second part (ADR normalization), the extracted ADR mentions are mapped to the standard concepts in MedDRA vocabulary. In the above example, the ADR mentions ‘*oww*’, ‘*toxic*’, and ‘*equivalent to torture*’ are mapped to the concepts ‘*pain (10033371)*’, ‘*drug toxicity (10013746)*’ and ‘*feeling unwell (10016370)*’ respectively. The dataset for this task consists of training set with 1862 tweets (1080 ADR tweets with 1464 ADR mentions and 782 nonADR tweets), validation set with 428 tweets (233 ADR tweets with 365 ADR mentions and 195 nonADR tweets) and test set with 976 tweets.

3.2 Methodology

ADR extraction is viewed as sequence labeling which is nothing but assigning a label to each of the tokens in the sequence. We follow BIO tagging: the tag ‘B-ADR’ represents tokens at the beginning of ADR mention, ‘I-ADR’ represents tokens inside ADR mention and ‘O’ represents nonADR tokens. We experiment with two models for this task. The first model is based on RoBERTa i.e., RoBERTaForTokenClassification. The second model is multi-task learning based RoBERTa. In this task, ADR tweet identification is the auxiliary task and ADR extraction is the main task. As these two tasks are similar in nature, by joint learning, the knowledge gained in auxiliary task improves the performance of the main task ADR extraction (Caruana, 1997; Crichton et al., 2017). Following the recent work in normalizing medical concepts (Kalyan and Sangeetha, 2020a; Subramanyam and Sivanesan, 2020), we treat concept normalization as multi-class classification and experiment with RoBERTa.

3.3 Experiments and Results

For ADR extraction, in case of RoBERTa based model, we set batch size = 64, epochs = 20 and learning rate = 0.00003. In case of multitask learning based RoBERTa model, $loss = \lambda \times L_{ADE} + (1 - \lambda) \times L_{ADR}$. Here L_{ADE} represents loss related to ADR extraction and L_{ADR} is loss related to ADR detection. The value of λ is set to 0.8. Here, the model is trained for 30 epochs with learning rate = 3e-5 and batch size = 64. For ADR normalization, we use a learning rate of 3e-5 and batch size of 128. For both ADR extraction and normalization, we use AdamW optimizer (Loshchilov and Hutter, 2019).

| Model | Evaluation | Type | Precision | Recall | F1 |
|----------------|------------|---------|-----------|--------|----------------------|
| RoBERTa | NER | Relaxed | 63.0 | 78.9 | 70.1 (↑ 13.7) |
| | | Strict | 41.1 | 54.2 | 46.8 |
| | NER + Norm | Relaxed | 30.4 | 39.8 | 34.5 |
| | | Strict | 23.6 | 31.1 | 26.8 |
| RoBERTa+MTL | NER | Relaxed | 65.1 | 72.8 | 68.7 |
| | | Strict | 45.2 | 52.5 | 48.6 |
| | NER + Norm | Relaxed | 32.6 | 37.7 | 35.0 (↑ 5.8) |
| | | Strict | 25.5 | 29.6 | 27.4 |
| Average scores | NER | Relaxed | 60.7 | 55.7 | 56.4 |
| | NER + Norm | Relaxed | 31.2 | 29.0 | 29.2 |

Table 2: Task 3 - ADR Extraction and Normalization results on test data. Here NER represents ADR Extraction and Norm represents ADR Normalization.

The performance of our models is reported in Table 2. From the table we observe that, a) RoBERTa based model achieved relaxed F1 score of 70.1% in ADR extraction which is 13.7% more than the relaxed average score b) Multi-task learning RoBERTa based model achieved relaxed F1 score of 35% in NER+Norm which is 5.8% more than the relaxed average score.

4 Conclusion

In this work, we explored the effectiveness of RoBERTa to identify, extract and normalize ADR mentions in tweets. In both task 2 and task 3, our proposed models achieved promising results with significant improvements over average scores.

References

- Rich Caruana. 1997. Multitask learning. *Machine learning*, 28(1):41–75.
- Gamal Crichton, Sampo Pyysalo, Billy Chiu, and Anna Korhonen. 2017. A neural network multi-task learning approach to biomedical named entity recognition. *BMC bioinformatics*, 18(1):368.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Katikapalli Subramanyam Kalyan and S Sangeetha. 2020a. Bertmcn: Mapping colloquial phrases to standard medical concepts using bert and highway network. Technical report, EasyChair.
- Katikapalli Subramanyam Kalyan and S Sangeetha. 2020b. Secnlp: A survey of embeddings in clinical natural language processing. *Journal of biomedical informatics*, 101:103323.
- Ari Z. Klein, Ilseyar Alimova, Ivan Flores, Arjun Magge, Zulfat Miftahutdinov, Anne-Lyse Minard, Karen O’Connor, Abeed Sarker, Elena Tutubalina, Davy Weissenbacher, and Graciela Gonzalez-Hernandez. 2020. Overview of the fifth social media mining for health applications (smm4h) shared tasks at coling 2020. In *Proceedings of the Fourth Social Media Mining for Health Applications (# SMM4H) Workshop & Shared Task*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Ilya Loshchilov and Frank Hutter. 2019. Decoupled weight decay regularization. In *International Conference on Learning Representations*.
- Xipeng Qiu, Tianxiang Sun, Yige Xu, Yunfan Shao, Ning Dai, and Xuanjing Huang. 2020. Pre-trained models for natural language processing: A survey. *arXiv preprint arXiv:2003.08271*.
- Abeed Sarker, Maksim Belousov, Jasper Friedrichs, Kai Hakala, Svetlana Kiritchenko, Farrokh Mehryary, Sifei Han, Tung Tran, Anthony Rios, Ramakanth Kavuluru, et al. 2018. Data and systems for medication-related text classification and concept normalization from twitter: insights from the social media mining for health (smm4h)-2017 shared task. *Journal of the American Medical Informatics Association*, 25(10):1274–1283.
- Kalyan Katikapalli Subramanyam and Sangeetha Sivanesan. 2020. Deep contextualized medical concept normalization in social media text. *Procedia Computer Science*, 171:1353 – 1362. Third International Conference on Computing and Network Communications (CoCoNet’19).
- Davy Weissenbacher, Abeed Sarker, Arjun Magge, Ashlynn Daughton, Karen O’Connor, Michael Paul, and Graciela Gonzalez. 2019. Overview of the fourth social media mining for health (smm4h) shared tasks at acl 2019. In *Proceedings of the Fourth Social Media Mining for Health Applications (# SMM4H) Workshop & Shared Task*, pages 21–30.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, pages arXiv–1910.