Adverse event recognition in Twitter: results from the WEB-RADR consortium

Lucie M. Gattepaille, Sara Hedfors Vidlin, Tomas Bergvall, Johan Ellenius. **Uppsala Monitoring Centre, Uppsala, Sweden**

Background

Pharmacovigilance consists in monitoring the safety of medicinal products throughout their life cycle. As clinical trials usually include a limited number of participants in a homogenous, non-fully representative population, constant post-marketing surveillance of medicines is an absolute necessity to detect rare and possibly serious adverse drug reactions (ADRs) as early as possible.

Traditionally, pharmacovigilance activities have relied on spontaneous reports describing cases of suspected ADRs. However, the rise and massive spread of social media over the last decade poses the question whether social media data can be harnessed and reliably utilized for pharmacovigilance purposes. Answering this question has been one of the aims of the WEB-RADR consortium, a large 3-year project involving partners from industry, regulatory agencies and academia, and supported by the Innovative Medicines Initiative [1].

With hundreds of millions of active users openly sharing their thoughts and experiences, Twitter has the potential to be a useful resource for pharmacovigilance, complementing by its unsolicited nature, timeliness and breadth of patient coverage the more traditional data sources [2]. We present the results of a pipeline of adverse event recognition in Twitter, developed in the course of the WED-RADR project.

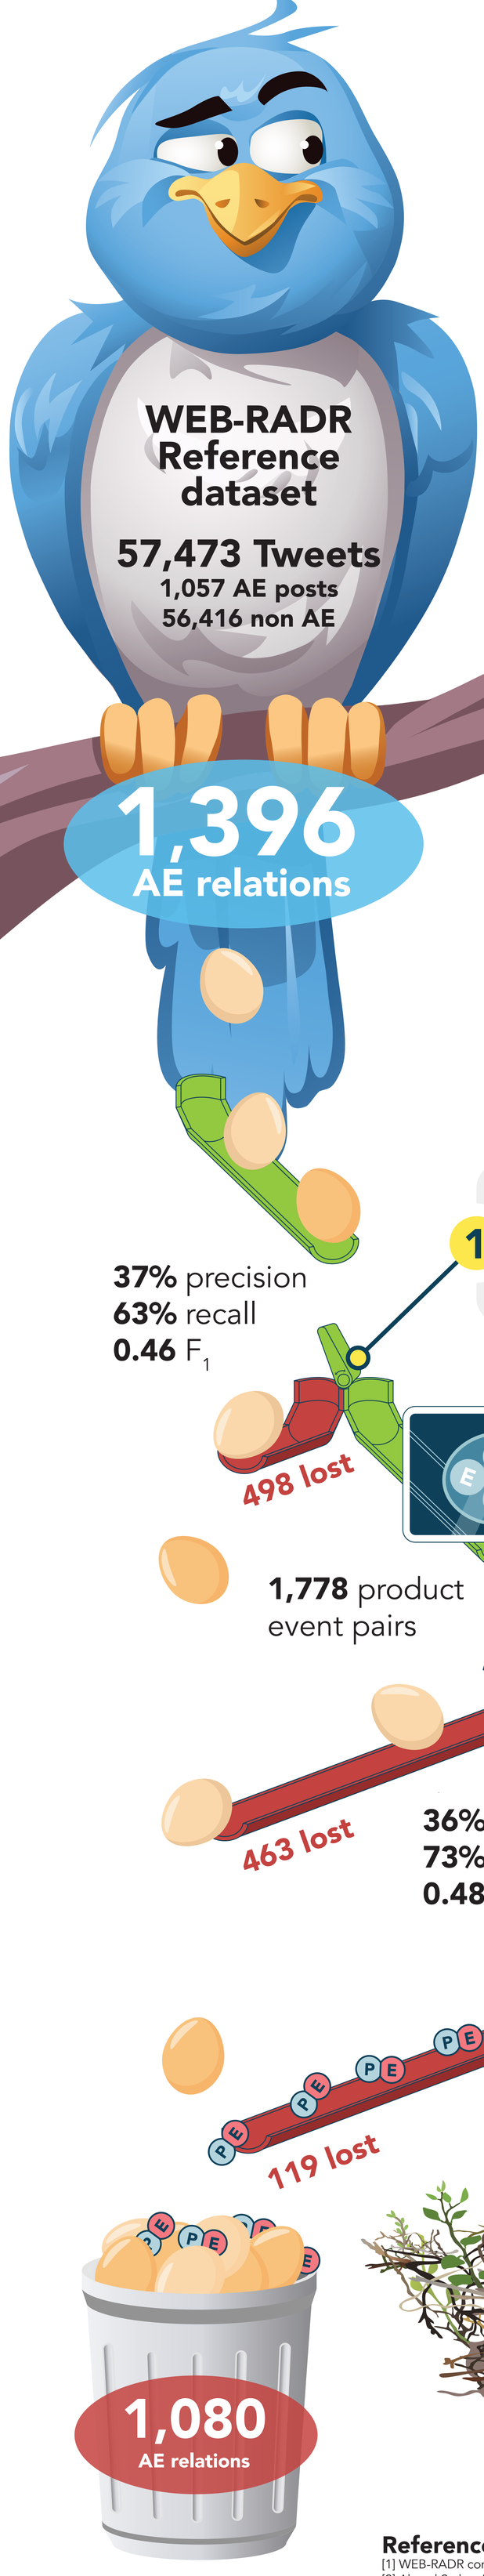

Aim

Implement a pipeline to recognize and normalize medicinal products (P) and medical events (E), and then characterize their relationship as an adverse event (AE) or not.

Methods

1. Remove tweets of low relevance
2. Perform Named Entity Recognition (NER) of medicinal product and medical event mentions. Normalize to well-established terminologies (WHODrug and MedDRA, the Medical Dictionary for Regulatory Activities)
3. Characterize every P/E pair according to the relationship implied by the author

Relevance filter

- Developed and published by Epidemico, a WEB-RADR collaborator [3]
- Score for resemblance to AE tweets
- Threshold 0.7

Datasets

For training the pipeline:
A proprietary manually-annotated dataset of around 145,000 Tweets, provided by WEB-RADR collaborators (training/validation/test split: 100,000/30,000/45,000).

For testing generalizability:
WEB-RADR Reference dataset of 57,473 manually annotated Tweets that will be made publicly available as a benchmark for comparison with future work on the task. The results provided here are given for this dataset.

NER module

P **Medicinal product** E **Medical event**

- Products recognized via dictionary lookup, filtered for ambiguous names (e.g. the product 'Today')
- Events recognized via dictionary lookups (MedDRA + vernacular extracted from VigiBase, the WHO global database of individual case safety reports) and logistic regressions using Tweets as bag-of-grams.

AE relation classifier

PE **Adverse event relation**

- Applied to every possible P/E pair in a Tweet
- Classifies whether the author implied an AE relation from the product to the medical event
- Logistic regression based on document features (e.g. number of URLs, of words, of user mentions), syntactic features (e.g. P before E, number of words between P and E) and semantic features using word2vec representations [4]

Results

Despite a good performance on the hold-out sample of the training dataset (96% recall of AE relations), the relevance filter provided by Epidemico poorly generalized on the WEB-RADR reference dataset (64% recall), leading to an early loss of 36% of the AE relations. Nonetheless, it is in the NER module that we proportionally lose most of the AE relations, by not managing to capture either the medicinal product or the medical event (48% recall, compared to 61% recall on the hold-out sample). Most of the true AE relations remaining when we apply the AE classifier do get recognized as AE relations (73% recall), but 64% of the suspected AE relations are false positives.

Overall, the pipeline achieved a precision of 36% and a recall of 23% (F1-score 0.28) on the WEB-RADR reference dataset.

In comparison

Results from an already published automated AE recognition pipeline [3].

	Published performance	Performance on the WEB-RADR Reference	Our performance
Precision	0.716	0.185	0.358
Recall	0.863	0.319	0.226
F_1	0.782	0.234	0.277

Discussion

This study highlights two major difficulties with developing methods of automatic detection of adverse events in Twitter posts: 1) detecting and normalizing medical events is a challenge in Twitter posts, probably due to the noisiness of the data (e.g. misspellings, abbreviations, diversity of layman expressions), 2) the transferability of models outside the universe of the training data to external datasets is poor, despite the use of a training/validation/test setup. The latter difficulty is poorly understood and should be the object of more research, to investigate the true ability of AE recognition algorithms to harness social media data.

References

[1] WEB-RADR consortium. https://web-radr.eu/about-us/, accessed May 7th, 2018
[2] Abeed Sarker, Rachel Ginn, Azadeh Nikfarjam, Karen O'Connor, Karen Smith, Swetha Jayaraman, Tejaswi Upadhaya, Graciela Gonzales. Utilizing social media data for pharmacovigilance: a review. Journal of Biomedical Informatics, 54:202-212, April 2015.
[3] Carrie E. Pierce, Khaled Bouri, Carol Pamer, Scott Proestel, Harold W. Rodriguez, Hoa Van Le, Clark C. Friefeld, John S Brownstein, Mark Walderhaug, I. Ralph Edwards, Nabarun Dasgupta. Evaluation of Facebook and Twitter monitoring to detect safety signals for medical products: an analysis of recent FDA safety alerts. Drug Safety, 40:317-334, April 2017.
[4] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S. Corrado, Jeff Dean. Distributed representations of words and phrases and their compositionality. In Proceedings of Advances in Neural Information Processing Systems, 26, 2013.

Uppsala Monitoring Centre

Uppsala Monitoring Centre (UMC)
Box 1051, SE-751 40 Uppsala, Sweden
+46 18 65 60 60 www.who-umc.org

"The research leading to these results was conducted as part of the WEB-RADR consortium, (http://webradr.eu) which is a public-private partnership coordinated by the Medicines and Healthcare products Regulatory Agency. The WEB-RADR project has received support from the Innovative Medicine Initiative Joint Undertaking (www.imi.europa.eu) under Grant Agreement n° 115632, resources of which are composed of financial contribution from the European Union's Seventh Framework Programme (FP7/2007-2013) and EFPIA companies' in kind contribution. The views expressed are those of the authors only."

WEB-RADR