Edinburgh Research Explorer

It is more than what you Say!

Link:
[Link to publication record in Edinburgh Research Explorer](#)

Document Version:
Peer reviewed version

It is more than what you Say!
Leveraging User Online Activity for Improved Stance Detection

Abeer ALDayel✷ Walid Magdy✤

Institute for Language, Computation and Cognition
School of Informatics, University of Edinburgh
`a.aldayel@ed.ac.uk`✷, `wmagdy@inf.ed.ac.uk`✤

Extended Abstract

To what extent do online interactions reveal an individual's viewpoint? Examining the implications of those interactions in detecting users' stances provides a better understanding of one of the pressing challenges in predicting viewpoints using only the raw texts of a user's posts. To address this challenge, we propose a new representation of stance detection that incorporates users' interaction elements to better predict their stances on a topic.

Stance detection is the task of inferring viewpoint towards a given topic or entity either being supportive or against (Mohammad et al., 2016). Lately, this problem has attracted considerable attention due to its value in studying rumours and fake news in social media (Ma et al., 2018). To date, most of the studies in this area have focused on using the textual elements of a user's posts independently from social factors, such as homophily and network structure. There has been some work on studying the integration of network and content with a limited focus on the ideological political views or specific event (Tsakalidis et al., 2018; Himelboim et al., 2013; Lai et al., 2016; Magdy et al., 2016). For instance the study of (Himelboim et al., 2013) focused on the liberal and conservative on twitter. Unlike previous work, rather than studying the stance on single topic and using a domain specific data, we study the stance in any domain. In this work, we analyse the effect of online social interactions on stance detection and demonstrate its superiority in predicting stances toward different topics. In addition, we show that complementing text disclosing stance with user's network interactions leads to significantly improved results.

We proposed two sets of network interactions to capture a better representation of a user's stance: (1) the activity network (NW_act), which includes the accounts and web domains with which the users directly interact with and cite in their own tweets, and (2) the preference network (NW_pref), which represents the indirect interactions with accounts and web domains in the posts the user likes. Using the preference network alleviates the problem of the silent (or cold) users and reveals the implicit stance expression towards a topic. Based on our suggested

Features	F$_{favour}$	F$_{against}$	F$_{avg}$
TXT	63.09	73.87	68.48
TXT_TK	68.07	82.09	75.08
NWact	64.04	76.18	70.11
NWpref	55.73	70.14	62.94
TXT+NWact	67.21	76.49	71.85
TXT_TK+NWact	71.25	82.42	**76.84**

Table 1: The result of two models SVM baseline *TXT* and Tree kernel SVM *TXT_TK* and when combining both text and network interactions representations.

set of features, which depends on the network interactions of a user's given posted tweet and is independent of the content of the tweet itself, it is expected that the stance-classifier model will always give the same classification to any tweets posted by that user. We argue that this is acceptable based on the assumption that a user's stance on a given topic is not expected to change within a short period of time (Borge-Holthoefer et al., 2015). Moreover, we hypothesize that, even if the user's stance does change, it would be accompanied by a change in the network interactions of the user (Waniek et al., 2018).

Our experiment was conducted with the popular benchmark of the SemEval-2016 stance-detection dataset (Mohammad et al., 2016). The dataset contains a set of 4000 tweets on five topics, including Atheism, Climate Change, the Feminist Movement, Hillary Clinton and the Legalisation of Abortion. To gauge the effectiveness of our new representation in predicting stances, we experimented with two models. The first model is the SVM model *TXT*, which achieved the top result among 19 teams in SemEval-2016 stance detection (Mohammad et al., 2016). The second model we implemented is the current state-of-the-art model, which achieved the top reported performance on SemEval-2016 stance data, *TXT_TK* and which used a tree representation of the tweets with a tree kernel SVM and ignored the none stance (Siddiqua et al., 2018). We compared the results of various combinations with a focus on gauging the improvements resulting from the network-interaction features (NW). Table 1 illustrates the overall results of exploiting the two types of network features along with content features using the two models. The results show that incorporating the (NW_act) with the content leads to a more accurate stance detection than using only content-based features. The obtained F-scores in this case are by far the highest results yet reported on the SemEval dataset, which confirms the significant impact of exploiting a user's network activity to boost the performance of stance detection, especially when combined with textual features.

Using network-interactions features without content helps in detecting the stance of users with implicit points of view toward a topic when they did not directly express their point of view by using keywords related to that topic. The method we used to detect the stance which exploits network features from two streams: the time-line stream and the favourites (likes) stream. Such features have no topical relation to the stance's target, yet the stance representation improved for the given topic. Our findings in this study show that, regardless of the topic, there are usually common signals in a user's activity and preference networks that indicate the user's stance towards the topic.

References

Borge-Holthoefer, J., Magdy, W., Darwish, K., and Weber, I. (2015). Content and network dynamics behind egyptian political polarization on twitter. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, pages 700–711. ACM.

Himelboim, I., McCreery, S., and Smith, M. (2013). Birds of a feather tweet together: Integrating network and content analyses to examine cross-ideology exposure on twitter. *Journal of Computer-Mediated Communication*, 18(2):154–174.

Lai, M., Farías, D. I. H., Patti, V., and Rosso, P. (2016). Friends and enemies of clinton and trump: using context for detecting stance in political tweets. In *Mexican International Conference on Artificial Intelligence*, pages 155–168. Springer.

Ma, J., Gao, W., and Wong, K.-F. (2018). Detect rumor and stance jointly by neural multi-task learning. In *Companion of the The Web Conference 2018*, pages 585–593.

Magdy, W., Darwish, K., Abokhodair, N., Rahimi, A., and Baldwin, T. (2016). # isisisnotislam or# deportallmuslims?: Predicting unspoken views. In *Proceedings of the 8th ACM Conference on Web Science*, pages 95–106. ACM.

Mohammad, S., Kiritchenko, S., Sobhani, P., Zhu, X.-D., and Cherry, C. (2016). SemEval-2016 Task 6: Detecting Stance in Tweets. In *SemEval@ NAACL-HLT*, pages 31–41.

Siddiqua, U. A., Chy, A. N., and Aono, M. (2018). Stance detection on microblog focusing on syntactic tree representation. In *International Conference on Data Mining and Big Data*, pages 478–490. Springer.

Tsakalidis, A., Aletras, N., Cristea, A. I., and Liakata, M. (2018). Nowcasting the stance of social media users in a sudden vote: The case of the greek referendum. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 367–376. ACM.

Waniek, M., Michalak, T. P., Wooldridge, M. J., and Rahwan, T. (2018). Hiding individuals and communities in a social network. *Nature Human Behaviour*, 2(2):139.