A Tutorial on Stance Detection

Dilek Küçük
TÜBİTAK MRC Energy Institute
Electrical Power Technologies Department
Ankara, Turkey
dilek.kucuk@tubitak.gov.tr

Fazli Can
Bilkent University
Computer Engineering Department
Ankara, Turkey
canf@cs.bilkent.edu.tr

ABSTRACT

Stance detection (also known as stance classification, stance prediction, and stance analysis) is a problem related to social media analysis, natural language processing, and information retrieval, which aims to determine the position of a person from a piece of text they produce, towards a target (a concept, idea, event, etc.) either explicitly specified in the text, or implied only. Common stance classes include *Favor*, *Against*, and *None*. In this tutorial, we will define the core concepts and other related research problems, present historical and contemporary approaches to stance detection (including shared tasks and tools employed), provide pointers to related datasets, and cover open research directions and application areas of stance detection. As solutions to stance detection can contribute to diverse applications including trend analysis, opinion surveys, user reviews, personalization, and predictions for referendums and elections, it will continue to stand as an important research problem, mostly on textual content currently, and particularly on Web content including social media.

CCS CONCEPTS

• **Computing methodologies** → **Natural language processing**; **Machine learning**; **Language resources**; • **Information systems** → **Information retrieval**; **Web and social media search**; **Sentiment analysis**.

KEYWORDS

Stance detection, Social media analysis, Twitter, Data streams

ACM Reference Format:
Dilek Küçük and Fazli Can. 2022. A Tutorial on Stance Detection. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining (WSDM '22), February 21–25, 2022, Tempe, AZ, USA.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3488560.3501391

1 INTRODUCTION

Stance detection has recently emerged as an important research problem in natural language processing (NLP), social media analysis, and information retrieval (IR), as revealed in the number of related papers published particularly since 2015 [15]. It is a research topic considered closely related to sentiment analysis and is commonly defined as the detection of the stance (as *Favor*, *Against*, or *None*) of the text producer towards a target [22, 23]. Apart from sentiment analysis, there are other research problems closely related to stance detection, including emotion detection, sarcasm detection, irony detection, controversy detection, and argument mining, among others.

Stance detection is known to have several application areas such as prediction of election/referendum results, information retrieval, rumour classification, and fake news detection. Particularly based on the latter two application areas, two subproblems of stance detection have emerged, namely, *rumour stance detection* and *fake news stance detection* [15]. Definitions of these subproblems along others (*multi-target stance detection* [27] and *cross-target stance detection* [39]) will be provided within the course of this tutorial.

An important contribution to the stance detection research is a series of shared tasks on stance detection between 2016 and 2017: an initial competition (within the course of SemEval-2016) on English tweets [22], another on Chinese microblog posts [40], and another related competition on Spanish and Catalan tweets [31]. More recently, a stance detection shared task has also been performed for tweets in Italian [6]. Another related shared task on Spanish and Basque tweets is again conducted recently [18]. In addition to the evaluation of different approaches for stance detection on microblogs in different languages, these competitions have facilitated the compilation significant annotated datasets [6, 18, 21, 22, 40] for stance detection.

Earlier work on stance detection utilized different rule-based and machine learning based algorithms on a variety of text genres including congressional-floor debates, online debate forums, student essays, and tweets [8, 9, 25, 32]. Related studies seem to have increased noticeably after the aforementioned initial shared task on stance detection in English tweets [22]. Hence, particularly considering those studies after (and including) this shared task; traditional machine learning approaches like SVM [1, 11, 12, 22, 23, 26, 38] and logistic regression [6, 24, 41, 42], deep learning approaches like LSTM [2, 7, 28, 29, 36] and CNN [35, 37, 43], and ensemble methods [20, 30, 33] have all been employed for the task of stance detection. Though machine learning approaches like SVM are the most commonly utilized ones up until 2019, more recent studies tend to apply deep learning algorithms, similar to the case of many tasks in NLP and IR [15]. There also exist significant research efforts to produce annotated datasets for stance detection [14, 17, 27, 44]. Currently, stance-annotated datasets exist for a number of languages including English [21, 27], Catalan [31], Chinese [40], Czech [10], Italian [19], Spanish [31], Turkish [13, 14], and Basque [18]. Yet, in addition to these language-specific datasets, recent work on the compilation of multilingual stance-annotated datasets [17, 34, 44] (where these datasets include annotated samples in other languages such as German and French) will facilitate related research and thereby will hopefully lead to important findings. There are several lines

WSDM '22, February 21–25, 2022, Tempe, AZ, USA

ACM ISBN 978-1-4503-9132-0/22/02…$15.00
https://doi.org/10.1145/3488560.3501391

of future research after having a firm base knowledge about the stance detection problem. Examples include cross-lingual and multilingual stance detection [15], stance detection in data streams [3], stance detection on non-textual content, context-sensitive stance detection [15], and exploring the application opportunities of new methods [4, 5].

This tutorial aims to cover the core concepts in stance detection together with related research topics, significant approaches to this problem published so far, relevant datasets and other practical resources, application areas, and finally, outstanding issues not sufficiently addressed in the related literature so far. We believe that this tutorial will be beneficial to interested researchers of Web search, NLP and IR alike, as well as related practitioners.

2 INTENDED AUDIENCE

Stance detection is an important research problem within the domains of social media analysis, Web search, natural language processing, and information retrieval. Therefore, we believe that researchers and practitioners of these research topics can readily benefit from this tutorial. There are not any particular prerequisites for the attendees.

3 OBJECTIVES OF THE TUTORIAL

The objective of this tutorial is to provide an overall understanding of the stance detection problem, which has significant and widespread application opportunities. The tutorial attendees will:

- grasp the main stance detection problem, its subproblems, and its relationship with other problems closely related to stance detection,
- learn about the evolution of the stance detection literature, and about common algorithms and approaches to this problem, in addition to the tools employed,
- be provided with pointers to related common stance detection datasets so that they can readily begin conducting stance detection experiments themselves on the available datasets, and
- learn about the related application areas as well as open research topics regarding stance detection, and learn about some common and favorable practices employed by stance detection researchers and practitioners.

4 PREVIOUS STANCE DETECTION TUTORIALS

To the best of our knowledge, there were two previous tutorials on stance detection. The first one was titled "*Detection and Characterization of Stance on Social Media*" which was carried out within the course of *14th International Conference on Web and Social Media (ICWSM-2020)*. There are several differences between our tutorial on stance detection and this previous tutorial: (1) Previous tutorial focuses on social media only while our tutorial will have a broader coverage including other input genres as in fake news stance detection where news articles are under consideration. (2) Being a more recent tutorial on the topic, our tutorial will cover more recent work which is a significant advantage as considerable body of work is being published on stance detection each year. (3) In our tutorial, we will pay particular attention to the subproblems of stance detection and other closely-related problems, in addition to significant application areas and outstanding issues.

The second tutorial on stance detection was our tutorial titled "*Stance Detection: Concepts, Approaches, Resources and Outstanding Issues*" which was carried out within the course of *44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR-2021)* [16]. Our current tutorial proposal will cover more recent work on stance detection, compared to the content of our previous tutorial. For instance, a new shared task named VaxxStance-2021 [18] has been performed in 2021 and work related to this recent shared task will be included in our tutorial. Similarly, tweet datasets on stance detection for new languages such as Basque [18] are being published within the course of more recent work. In this tutorial at SIGIR-2021, there were questions from the audience regarding the differences and interrelationships between stance detection and related problems such as sentiment analysis and controversy detection. We addressed these questions according to the findings of the related work [15, 23]. There were also questions regarding the annotation procedures of stance detection datasets and these questions were also answered based on the related previous work on stance detection and related problems. The audience was provided with pointers to relevant literature on stance detection.

5 TUTORIAL OUTLINE

The outline of our tutorial on stance detection is presented below.

(1) Introduction
(2) Core concepts and related problems
    - Definitions of stance detection and subproblems of stance detection.
    - Interrelationships with related problems including sentiment analysis, emotion recognition, sarcasm/irony detection, controversy detection, and argument mining
(3) Stance detection competitions (shared tasks)
    - SemEval-2016 Task 6: Detecting Stance in Tweets
    - Shared Task of Stance Detection in Chinese Microblogs at NLPCC-ICCPOL-2016
    - Shared Task of Stance Detection in Spanish and Catalan Tweets at IberEval-2017
    - SardiStance: Stance Detection Task in Italian Tweets at EVALITA-2020
    - VaxxStance: Going Beyond Text in Cross-lingual Stance Detection at IberLeF-2021
(4) Historical and contemporary approaches
    - Earlier studies on stance detection
    - Machine learning based approaches
    - Deep learning based approaches
    - Ensemble learning approaches
    - Related learning systems and tools
(5) Common stance detection datasets
(6) Application areas
    - Opinion surveys/polling
    - Public health surveillance
    - Information retrieval
    - Stance summarization
    - Rumour classification

- Fake news detection

(7) Outstanding issues
- Stance detection in data streams
- Cross-lingual and multilingual stance detection
- Context-sensitive stance detection

(8) Conclusion

REFERENCES

[1] Aseel Addawood, Jodi Schneider, and Masooda Bashir. 2017. Stance classification of Twitter debates: the encryption debate as a use case. In *Proceedings of the 8th International Conference on Social Media & Society*. 2.

[2] Rabab Alkhalifa and Arkaitz Zubiaga. 2020. QMUL-SDS at SardiStance2020: Leveraging network interactions to boost performance on stance detection using knowledge graphs. *arXiv preprint arXiv:2011.01181* (2020).

[3] Hamed R Bonab and Fazli Can. 2018. GOOWE: geometrically optimum and online-weighted ensemble classifier for evolving data streams. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 12, 2 (2018), 25.

[4] Hamed R Bonab and Fazli Can. 2019. Less is more: a comprehensive framework for the number of components of ensemble classifiers. *IEEE Transactions on Neural Networks and Learning Systems* 30, 9 (2019), 2735–2745.

[5] Umit Can and Bilal Alatas. 2021. A novel approach for efficient stance detection in online social networks with metaheuristic optimization. *Technology in Society* 64 (2021), 101501.

[6] Alessandra Teresa Cignarella, Mirko Lai, Cristina Bosco, Viviana Patti, Rosso Paolo, et al. 2020. SardiStance@EVALITA2020: Overview of the task on stance detection in Italian tweets. In *EVALITA 2020 Seventh Evaluation Campaign of Natural Language Processing and Speech Tools for Italian*. 1–10.

[7] Kuntal Dey, Ritvik Shrivastava, and Saroj Kaushik. 2018. Topical stance detection for Twitter: A two-phase LSTM model using attention. In *European Conference on Information Retrieval*. 529–536.

[8] Adam Faulkner. 2014. Automated classification of stance in student essays: an approach using stance target information and the Wikipedia link-based measure. In *Proceedings of the International Florida Artificial Intelligence Research Society Conference*. 174–179.

[9] Kazi Saidul Hasan and Vincent Ng. 2013. Stance classification of ideological debates: data, models, features, and constraints. In *Proceedings of the International Joint Conference on Natural Language Processing*. 1348–1356.

[10] Tomáš Hercig, Peter Krejzl, Barbora Hourová, Josef Steinberger, and Ladislav Lenc. 2017. Detecting stance in Czech news commentaries. In *Proceedings of the Conference on Theory and Practice of Information Technologies (ITAT)*.

[11] Dilek Küçük. 2017. Stance detection in Turkish tweets. In *Proceedings of the International Workshop on Social Media World Sensors*.

[12] Dilek Küçük and Fazli Can. 2018. Stance detection on tweets: an SVM-based approach. *arXiv preprint arXiv:1803.08910* (2018).

[13] Dilek Küçük. 2021. Sentiment, stance, and intent detection in Turkish tweets. In *New Opportunities for Sentiment Analysis and Information Processing*. IGI Global, 206–217.

[14] Dilek Küçük and Fazli Can. 2019. A tweet dataset annotated for named entity recognition and stance detection. *arXiv preprint arXiv:1901.04787* (2019).

[15] Dilek Küçük and Fazli Can. 2020. Stance detection: A survey. *ACM Computing Surveys (CSUR)* 53, 1 (2020), 1–37. https://doi.org/10.1145/3369026

[16] Dilek Küçük and Fazli Can. 2021. Stance detection: Concepts, approaches, resources, and outstanding issues. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2673–2676.

[17] Mirko Lai, Alessandra Teresa Cignarella, Delia Irazú Hernández Farías, Cristina Bosco, Viviana Patti, and Paolo Rosso. 2020. Multilingual stance detection in social media political debates. *Computer Speech & Language* (2020), 101075.

[18] Mirko Lai, Alessandra Teresa Cignarella, Finos Livio, and Andrea Sciandra. 2021. WordUp! at VaxxStance 2021: Combining contextual information with textual and dependency-based syntactic features for stance detection. In *Proceedings of the International Conference of the Spanish Society for Natural Language Processing*.

[19] Mirko Lai, Viviana Patti, Giancarlo Ruffo, and Paolo Rosso. 2018. Stance evolution and Twitter interactions in an Italian political debate. In *Proceedings of 23rd International Conference on Natural Language and Information Systems*.

[20] Can Liu, Wen Li, Bradford Demarest, Yue Chen, Sara Couture, Daniel Dakota, Nikita Haduong, Noah Kaufman, Andrew Lamont, Manan Pancholi, Kenneth Steimel, and Sandra Kübler. 2016. IUCL at SemEval-2016 task 6: an ensemble model for stance detection in Twitter. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*. 394–400.

[21] Saif M Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. A dataset for detecting stance in tweets. In *Proceedings of the Language Resources and Evaluation Conference*. 3945–3952.

[22] Saif M Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. SemEval-2016 task 6: detecting stance in tweets. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*. 31–41.

[23] Saif M Mohammad, Parinaz Sobhani, and Svetlana Kiritchenko. 2017. Stance and sentiment in tweets. *ACM Transactions on Internet Technology* 17, 3 (2017), Article 26.

[24] Mauridhi Hery Purnomo, Surya Sumpeno, Esther Irawati Setiawan, and Diana Purwitasari. 2017. Biomedical engineering research in the social network analysis era: stance classification for analysis of hoax medical news in social media. *Procedia Computer Science* 116 (2017), 3–9.

[25] Ashwin Rajadesingan and Huan Liu. 2014. Identifying users with opposing opinions in Twitter debates. In *Proceedings of the International Conference on Social Computing, Behavioral-Cultural Modeling, and Prediction*. 153–160.

[26] Anirban Sen, Manjira Sinha, Sandya Mannarswamy, and Shourya Roy. 2018. Stance classification of multi-perspective consumer health information. In *Proceedings of the ACM India Joint International Conference on Data Science and Management of Data*. 273–281.

[27] Parinaz Sobhani, Diana Inkpen, and Xiaodan Zhu. 2017. A dataset for multi-target stance detection. In *Proceedings of the Conference of the European Chapter of the Association for Computational Linguistics*. 551–557.

[28] Qingying Sun, Zhongqing Wang, Shoushan Li, Qiaoming Zhu, and Guodong Zhou. 2018. Stance detection via sentiment information and neural network model. *Proceedings of the Frontiers of Computer Science* (2018).

[29] Qingying Sun, Zhongqing Wang, Qiaoming Zhu, and Guodong Zhou. 2018. Stance detection with hierarchical attention network. In *Proceedings of the International Conference on Computational Linguistics*. 2399–2409.

[30] Sahil Swami, Ankush Khandelwal, Vinay Singh, Syed Sarfaraz Akhtar, and Manish Shrivastava. 2018. An English-Hindi code-mixed corpus: stance annotation and baseline system. *arXiv preprint arXiv:1805.11868* (2018).

[31] Mariona Taulé, M Antonia Martí, Francisco Rangel, Paolo Rosso, Cristina Bosco, and Viviana Patti. 2017. Overview of the task on stance and gender detection in tweets on Catalan independence at IberEval 2017. In *Proceedings of the Second Workshop on Evaluation of Human Language Technologies for Iberian Languages (IberEval 2017)*.

[32] Matt Thomas, Bo Pang, and Lillian Lee. 2006. Get out the vote: determining support or opposition from congressional floor-debate transcripts. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. 327–335.

[33] Martin Tutek, Ivan Sekulic, Paula Gombar, Ivan Paljak, Filip Culinovic, Filip Boltuzic, Mladen Karan, Domagoj Alagić, and Jan Šnajder. 2016. Takelab at SemEval-2016 task 6: stance classification in tweets using a genetic algorithm based ensemble. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*. 464–468.

[34] Jannis Vamvas and Rico Sennrich. 2020. X-stance: A multilingual multi-target dataset for stance detection. *arXiv preprint arXiv:2003.08385* (2020).

[35] Prashanth Vijayaraghavan, Ivan Sysoev, Soroush Vosoughi, and Deb Roy. 2016. DeepStance at SemEval-2016 task 6: detecting stance in tweets using character and word-level CNNs. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*.

[36] Penghui Wei, Junjie Lin, and Wenji Mao. 2018. Multi-target stance detection via a dynamic memory-augmented network. In *Proceedings of the International ACM SIGIR Conference on Research & Development in Information Retrieval*. 1229–1232.

[37] Wan Wei, Xiao Zhang, Xuqin Liu, Wei Chen, and Tengjiao Wang. 2016. pkudblab at SemEval-2016 task 6: a specific convolutional neural network system for effective stance detection. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*. 384–388.

[38] Michael Wojatzki and Torsten Zesch. 2016. ltl.uni-due at SemEval-2016 task 6: stance detection in social media using stacked classifiers. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*. 428–433.

[39] Chang Xu, Cecile Paris, Surya Nepal, and Ross Sparks. 2018. Cross-target stance classification with self-attention networks. *arXiv preprint arXiv:1805.06593* (2018).

[40] Ruifeng Xu, Yu Zhou, Dongyin Wu, Lin Gui, Jiachen Du, and Yun Xue. 2016. Overview of NLPCC shared task 4: stance detection in Chinese microblogs. In *Natural Language Understanding and Intelligent Applications*. 907–916.

[41] Shaodian Zhang, Lin Qiu, Frank Chen, Weinan Zhang, Yong Yu, and Noémie Elhadad. 2017. We make choices we think are going to save us: debate and stance identification for online breast cancer CAM discussions. In *Proceedings of the International Conference on World Wide Web Companion*. 1073–1081.

[42] Zhihua Zhang and Man Lan. 2016. ECNU at SemEval 2016 task 6: relevant or not? supportive or not? a two-step learning system for automatic detecting stance in tweets. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*. 451–457.

[43] Yiwei Zhou, Alexandra I Cristea, and Lei Shi. 2017. Connecting targets to tweets: semantic attention-based model for target-specific stance detection. In *Proceedings of the International Conference on Web Information Systems Engineering*. 18–32.

[44] Elena Zotova, Rodrigo Agerri, and German Rigau. 2021. Semi-automatic generation of multilingual datasets for stance detection in Twitter. *Expert Systems with Applications* 170 (2021), 114547.