Stance Trees: A Novel Approach for Assessing Politically Polarized Issues in Twitter

Gabriela Andrea Díaz Carlos Iván Chesñevar Elsa Estevez Ana Maguitman
Departamento de Ciencias e Ingeniería de la Computación,
Universidad Nacional del Sur
Instituto de Cs. e Ing. de la Computación (ICIC UNS–CONICET)
Bahía Blanca, Buenos Aires, Argentina
{gabriela.diaz,cic,ece,agm}@cs.uns.edu.ar

ABSTRACT

Social and political polarization, which sometimes is the result of misinformation, is a common obstacle that can be harmful at the moment of communicating government policies. Intelligent tools that aid critical thinking in the light of different opinions and standpoints available in social media can help ameliorate this obstacle. This paper presents preliminary research work toward developing such tools by proposing a methodology for building stance trees based on tweets collected from social media. Stance trees are hierarchical structures where nodes represent arguments pro, anti, or uncertain about a target issue and edges stand for attack relations between those arguments. The proposed methodology includes retrieving tweets relevant to the target issue, manually labeling a sample set of the collected tweets, developing and applying a model for stance detection, and finally building a stance tree. We illustrate the expected results through a case study on the politically polarized "COVID-19 vaccine" issue. Our preliminary results demonstrate the feasibility of the proposal and highlight the utility of stance trees as a tool for aiding critical thinking.

CCS CONCEPTS

• **Information systems** → **Data mining**; • **Applied computing** → **E-government**; • **Computing methodologies** → *Supervised learning by classification.*

KEYWORDS

Artificial Intelligence, E-Participation, Stance Detection, Social Media, Argumentation, Political Polarization

ACM Reference Format:
Gabriela Andrea Díaz Carlos Iván Chesñevar Elsa Estevez Ana Maguitman. 2022. Stance Trees: A Novel Approach for Assessing Politically Polarized Issues in Twitter. In *15th International Conference on Theory and Practice of Electronic Governance (ICEGOV 2022), October 4–7, 2022, Guimarães, Portugal.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3560107.3560296

ICEGOV 2022, October 4–7, 2022, Guimarães, Portugal

ACM ISBN 978-1-4503-9635-6/22/10…$15.00
https://doi.org/10.1145/3560107.3560296

1 INTRODUCTION

Social media typically reflects citizens' concerns about various issues, giving rise to different standpoints and controversial viewpoints. Some of such issues may be the direct result of certain government policies (e.g. "COVID-19 vaccine", "legalization of abortion", "right to bear arms", and "feminism", among many others), engendering controversies which might result in a strong social and political polarization. While disagreement and argumentation are healthy and natural in any society, polarization may be harmful since it may lead to segregation, antagonism, deception, violence, and losing trust in key institutions [19]. Sometimes polarization is the result of misinformation, which is a common obstacle that can greatly mislead a significant part of the society. There is probably no definitive solution that a government can adopt to fight the problems of polarization and misinformation. Clearly, adopting policies for controlling speech or disabling the flow of information are far from being acceptable approaches, as they may jeopardize democratic free speech and promote even more polarization and intolerance. However, intelligent tools that aid critical thinking in the light of different opinions and standpoints available in social media can play an important role in improving social awareness. Developing such tools requires implementing effective methods for *stance detection* in social media.

Stance detection is the task of automatically determining from text whether the author of the text is in favor of, against, or neutral towards a proposition or target [2, 25]. Stance detection is an opinion mining task that is different from sentiment analysis in that the former focuses on the standpoint towards a given proposition while the latter looks into the emotional tone behind a body of text. Although the identification of sentiments from social media content has been extensively analyzed [4, 11, 18, 26, 32], less work has been carried out on the problem of stance detection [2, 12, 20, 29].

This work proposes to apply stance detection to a set of tweets about an issue under analysis to build a **stance tree**. A stance tree is a hierarchical structure that can be constructed by analyzing the prevailing stance (pro, anti or uncertain) associated with different aspects of the issue under analysis. Along with that approach, argumentation comes into play as a way of contrasting alternative standpoints and facilitating critical thinking about the issue at hand. In this particular approach, arguments in a stance tree are a cohesive set of tweets referring to a particular issue, whereas the structural relationship among arguments is given by indicating attacks (which are in conflict with a particular argument at issue). The dialogical and assessment layers provide a global view of all the arguments involved in the process.

In recent previous work, opinion trees [5, 14] were proposed as a way of contrasting collective opinions. However, opinion trees are based on sentiment analysis, focusing on providing a prevailing sentiment (positive, negative or neutral) associated with a set of tweets. In that respect, the concept of stance tree represents a natural evolution of the concept of opinion tree, as the notion of stance can be more naturally associated with an argument that is part of a dialectical process.

In this paper, we present a methodology to build a stance tree from a set of tweets. Based on a case study about the "COVID-2019 vaccine" issue we illustrate the proposed approach. We also discuss the applicability of our proposal in the context of identifying misinformation, helping thus to intelligent decision making from government stakeholders in the face of polarized issues.

2 RELATED WORK

The analysis of political polarization on social media has a long tradition and has been mostly studied by analyzing the network topologies induced by mentions, retweets or followers-followees relations [1, 7, 31] On the other hand, research on the impact of misinformation is relatively a new research topic, associated from the very beginning with politically polarized issues such as elections and crisis management [9, 21]. In the last two years, the COVID-19 pandemic and vaccination strategies worldwide [30] as well as the last US presidential election [27] were the two major events which deserved particular attention in stance detection and the impact of misinformation in Twitter.

Recent research on stance detection in Twitter has a statistical background, providing different metrics for handling misinformation and polarized opinions (e.g. [22]). The COVID-19 outbreak has also prompted the development of intelligent systems based on machine learning techniques for automatically identifying relevant tweets deemed to be informative, as done in [24]. In contrast with our proposal, this system provides a classification for removing tweets that are deemed 'uninformative' for use in down- stream epidemiological analyses, without providing a more structured view or taxonomy. In [30], the authors rely on a neural framework for the identification of stance towards misinformation about COVID-19 vaccines. Their approach is based on knowledge graphs, aiming at stance identification using the so-called attitude consistency (AC), modeled through a knowledge graph. Our proposal provides an alternative structure for stance detection, aiming at an easier identification of key terms in handling politically polarized issues, and providing a more natural way of representing different alternative arguments (pro, anti and neutral).

In [3], the authors propose three effective misinformation detection models for Twitter in the context of COVID-19, based on different machine learning techniques (long short-term memory (LSTM) networks, multichannel convolutional neural network (MC-CNN); and k-nearest neighbors (KNN)). Simulations are conducted to evaluate the performance of the proposed models in terms of various evaluation metrics. The focus of this research is based on assessing performance evaluation metrics when identifying individual tweets, rather than identifying the underlying arguments associated with a set of tweets and their interrelationships, as done in our approach.

In [5], the authors advocate for the use of computational argumentation as an underlying layer for modeling intelligent decision making in cognitive cities, highlighting the role of "opinion trees" [13, 14] when handling conflicting opinions in social media. They contend that the information made available in Twitter can be useful to extract a particular version of computational arguments (called "opinions") which emerge bottom-up from the social interaction associated with such messages. Opinions can be thus mined from Twitter based on incrementally generated queries, resulting in an "opinion tree", rooted in the first original query. Our approach based on "stance trees" borrows this key concept for computational argumentation (a tree-like representation for contrasting potentially contradictory opinions), but the process for obtaining such a tree is significantly different. In [14], an argument is given by a set of tweets syntactically associated with a query Q, whereas in our approach we generate arguments by computing a cohesive set of tweets using a soft clustering algorithm. Thus, our approach can be seen as a natural extension of opinion trees, relying on a supervised model for identifying polarization (rather than just a syntactic search on the Twitter database).

3 PROPOSED METHODOLOGY

The process of building a stance tree for a target issue using social media data involves multiple steps. A schematic overview of the proposed methodology is presented in Figure 1. Initially, a set of tweets relevant to the issue under analysis is collected (step 1). This can be done by implementing a topic-based retrieval tool for Twitter or by simply using an existing dataset with tweets related to the issue of interest.

Building a supervised model for stance detection requires accessing a labeled dataset. A random sample dataset is selected and each tweet from the sample dataset is labeled according to its stance towards the target issue as pro, anti or neutral (step 2). Following the usual machine learning methodology, the labeled dataset is used to train, validate and test a model (step 3). While any multi-class classifier can be used for this purpose, state-of-the-art BERT-based transformer networks [10] are an effective choice that can be fine-tuned to build a prediction model for the problem at hand. The resulting model can then be applied for stance detection on tweets relevant to the target issue under analysis (step 4). This way, a predicted stance (pro, anti or neutral) can be associated with new tweets. Neutral tweets are discarded as they are not useful for arguing in favor or against a position. Finally, a stance detection tree is built by applying a novel algorithm adapted from the one used to build opinion trees [14] (step 5).

Similar to opinion trees, stance trees resemble a dialectical tree where each node in the tree represents an argument that can take three possible polarities. However, the process required for building a stance tree is significantly different from the process of building an opinion tree. Each node in a stance tree represents an argument A_i based on a cohesive set of tweets T_i. The cohesive sets of tweets are generated by applying a soft clustering algorithm to the tweets dataset. Since soft clustering algorithms can assign the same object to more than one cluster it is possible that $T_i \cap T_j \neq \emptyset$ for $i \neq j$ when there exist tweets in T_i and T_j that share a common theme.

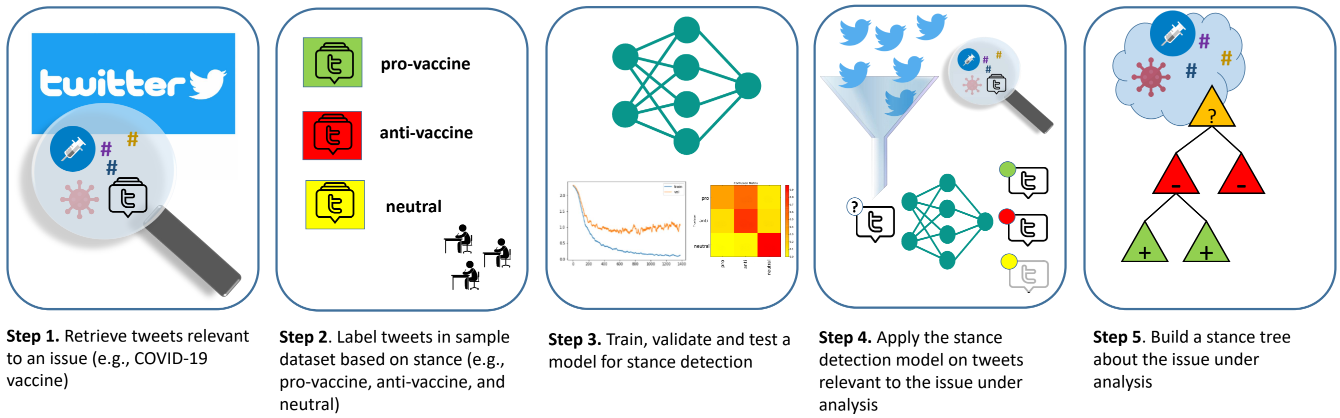

Figure 1: Proposed methodology for building a stance tree on an issue of interest. The pipeline for building a stance tree from social media involves (1) retrieving tweets relevant to an issue of interest using a topic-based information retrieval approach, (2) creating an annotated dataset with tweets labeled as pro, anti or neutral with respect to the issue at hand, (3) training, validating and testing a stance detection model, (4) applying the stance detection model to new tweets relevant to the issue under analysis, and (5) building the stance tree about the issue under analysis.

We associate with each argument A_i in a stance tree the prevailing stance of the tweets in T_i, which will typically be pro (labeled as "+") or anti (labeled as "-"). However, if T_i is highly polarized, the stance of A_i will be uncertain (labeled as "**?**"). Based on the prevailing stance associated with each argument, we say that argument A_i attacks argument A_j if the following three conditions hold: (1) A_i and A_j have different polarities (i.e., they are conflicting arguments), (2) $T_i \cap T_j \neq \emptyset$ (i.e., some tweets in different clusters share a common theme), and (3) the prevailing polarity of $T_i \cap T_j$ is equal to the prevailing polarity of T_i (i.e., the polarity of those tweets in different clusters that share a common theme tends to be equal to the polarity of A_i).

To illustrate the idea behind these three conditions consider the highly polarized "COVID-19 vaccine" issue. A possible argument A_i may be based on a set of tweets that refer to "depopulation", which is a common reference in anti-vaccine tweets. On the other hand, a pro-vaccine argument A_j may be based on tweets that talk about "COVID-19 eradication". Note that for this example condition (1) is satisfied (A_i and A_j have different polarity). Also, condition (2) is satisfied if there are tweets from both clusters that share a common theme ($T_i \cap T_j \neq \emptyset$). Finally, if the prevailing polarity of those tweets from both clusters with a common themes is negative then condition (3) is satisfied (the prevailing polarity of $T_i \cap T_j$ is equal to the prevailing polarity of T_i). In this case, we can conclude that the "depopulation"-based argument attacks the "COVID-19 eradication"-based argument. Next section presents a case study that offers an overview of how a stance tree for the "COVID-19 vaccine" issue can be built by applying the proposed methodology.

4 CASE STUDY

For this case study, we use the ANTi-Vax dataset [15], which comprises over 15,000 tweets related to the "-19 vaccine" topic collected from November 2020 to July 2021. The tweets in this dataset are labeled as misinformation (1) or no-misinformation (0). However, this dataset does not contain any stance information.

Once the tweets were hydrated the dataset was reduced to 12,228 tweets. A random sample of 1,050 tweets was selected and manually classified by three annotators into three categories: pro-vaccine, anti-vaccine and neutral. A few examples of tweets labeled into each category are presented in table 1.

The annotation process was completed in two stages. Initially, a common subset of 50 tweets was labeled by the three annotators to be able to calculate the inter-annotator agreement. The Cohen's kappa coefficient value [6] averaged across the three combinations of annotator pairs is 0.84, which represents an almost perfect agreement. In light of the high inter-annotator agreement and to expedite the process of labeling the remaining 1,000 tweets, during the second annotation stage each annotator was assigned a disjoint set of tweets of similar size. Using the labeled dataset we trained a BERT-based transformer network [10] and achieved ~85% accuracy (details on the evaluations will be presented in a future paper).

To present an illustrative example of the expected results, we manually analyzed the dataset of tweets to identify potential arguments associated with the "COVID-19 vaccine" issue. A stance tree obtained from this analysis is presented in Figure 2. In this example, it is possible to see that the stance about "COVID-19 vaccine" is highly polarized and therefore it is uncertain (marked with "?" on the root node of the stance tree). However, when different aspects associated with this issue are analyzed, different pro and anti standpoints arise. In particular, several anti-vaccine tweets refer to the fact that "more tests are needed" to guarantee safety. However, some pro-vaccine tweets state that vaccines "save lives". Others claim that vaccines will bring back "normality". Some polarization arises concerning going "back to work", with a pro argument associated with measures to "reactivate economy". The discussion could go further, including claims such as the fact that "normality" results in more "pollution", etc.

pro	**P1.** I am now half vaccinated... a moderna man
	P2. There's a new antivaccine lie, the claim that an mRNA #CovidVaccine (e.g., like the ones from @pfizer and @moderna_tx) is not a "vaccine" but a "medical device" or "gene therapy" and was falsely classified as a vaccine in order to bypass safety testing ...
	P3. Do you know what eradicated the most deadly illness from the face of the earth? Vaccination did!
anti	**A1.** Pfizer admitting that those injected with the vaccine are shedding. Environmental exposure they call it. very worrying stuff in there about infertility. If you've had this jab in the last 45 days as that's what the paper itself suggests please stay away fro
	A2. My mom and all my aunts & uncles, who supposed to take the jab, all vigorously refuse to take the experimental, gene altering, magnetic, nano-particle laced vaccine!!!
	A3. I hate to see this coming to pass, because too many of my friends, families and their children have had the experimental gene therapy.
neutral	**N1.** people asking which vaccine you got is the same vibe as asking which team you were in pokemon go
	N2. can someone explain what people mean on titkok when they're saying "if you kin (person) then you should/shouldn't worry about what's in the vaccine" like ik the rona vaccine but what do they mean by that

Table 1: Illustrative examples of tweets labeled as pro-vaccine, anti-vaccine and neutral. By analyzing these examples, it is possible to identify some challenging situations. For instance, tweet P2 and A3 use similar vocabulary (e.g., "gene therapy"). However, the former is pro-vaccine while the latter is anti-vaccine.

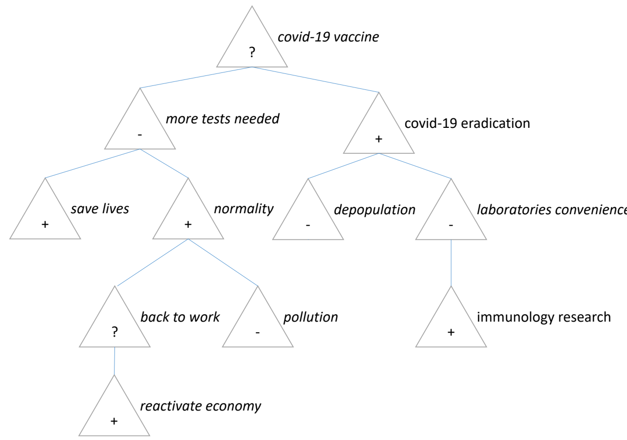

Figure 2: Illustrative example of a stance tree on the "covid-19 vaccine" issue. The issue "covid-19 vaccine" is highly polarized. There are two main arguments, namely an anti-vaccine argument stating that "more tests are needed" and a pro-vaccine argument claiming that the vaccine favors "covid-19 eradication". The resulting stance tree provides a global view of several pro- and anti-vaccine arguments.

	anti	pro	neutral
no-misinformation	17	612	40
misinformation	308	11	12

Table 2: Correlation matrix for misinformation vs. stance. We observe that pro-vaccine is highly correlated with no-misinformation, while anti-vaccine is highly correlated with misinformation.

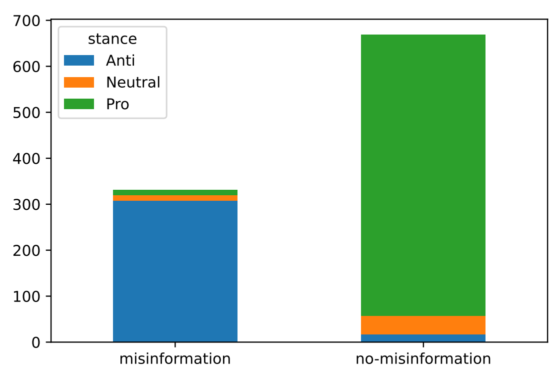

Figure 3: Stance vs. misinformation. We observe that most tweets labeled as pro-vaccine are classified as no-misinformation, while most tweets labeled as anti-vaccine are labeled as misinformation.

An interesting question that arises is whether there is a correspondence between stance and misinformation labels. To this end, we contrasted the tweets with stance and misinformation labels from the ANTi-Vax dataset [15]. This analysis is presented in the bar chart of Figure 3 and correlation matrix of Table 2, where we can observe that most tweets labeled as pro-vaccine are classified as no-misinformation. Correspondingly, most tweets labeled as anti-vaccine, are labeled as misinformation. This case study highlights the utility of stance trees as a tool for aiding critical thinking and potentially fighting misinformation in the face of polarized issues. Finally, to gain preliminary insight on the evolution of the stance towards the "covid-19 vaccine" issue, we present in Figure 4 a bar chart showing the number of tweets for each position (anti, neutral and pro) identified in the labeled set during the period January 2021-July 2021. We observe that the number of pro-vaccine tweets significantly increases in April 2021 and the number of anti-vaccine tweets is higher than the number of pro-vaccine ones in May 2021. A month-by-month chart showing the relation between stance and misinformation is presented in Figure 5. Similar to what we observe in Figure 3, this monthly analysis shows that most tweets labeled as pro-vaccine are classified as no-misinformation and most tweets labeled as anti-vaccine are classified as misinformation.

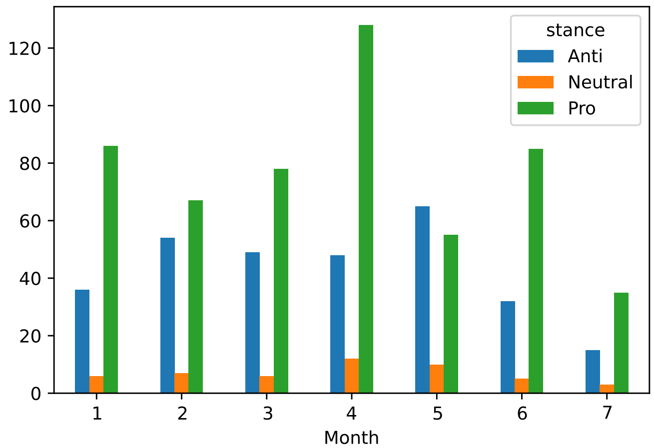

Figure 4: Stance evolution (anti, neutral and pro) during the period January 2021 - July 2021. We observe a significant increase in the number of pro-vaccine tweets in April 2021. We also observe that the number of anti-vaccine tweets is higher than the number of pro-vaccine ones in May 2021

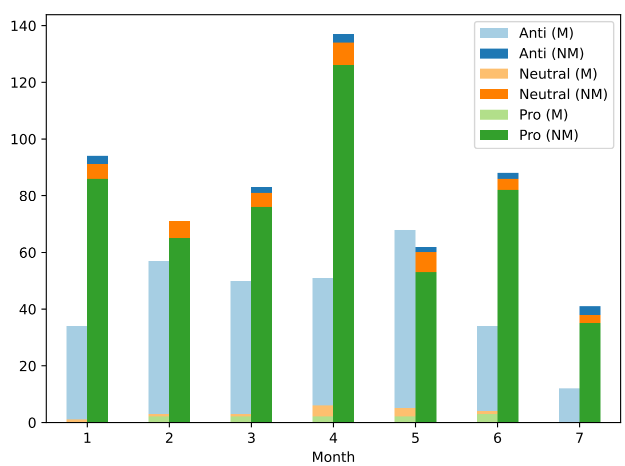

Figure 5: Stance (anti, neutral and pro) vs. misinformation during the period January 2021 - July 2021 analyzed on a month-by-month basis.

5 CONCLUSION

In this paper we have presented a novel approach for critically assessing politically polarized issues on Twitter. The proposed approach involves training a stance detection model on an issue under analysis by applying a traditional machine learning methodology. Subsequently, the trained model is used for stance detection on a set of tweets. A soft clustering algorithm is applied to group tweets that share a common theme. Each resulting cluster allows identifying arguments about the issue at hand. Finally, we formulated a definition of argument attack and outlined a method for computing attack relations on tweets sets. The proposed approach allows building a stance tree on the issue under analysis, which is a hierarchical structure where nodes represent arguments pro, anti, or uncertain about a target issue and edges stand for attack relations between those arguments. A case study on the "COVID-19 vaccine" issue illustrates the practical aspects of the approach.

We contend that our approach can be thought of as an example of *argumentative* Explainable AI [8], where computational argumentation provides the backbone for a wide array of reasoning abstractions and explanation delivery methodologies. In this way, stance trees can naturally complement other existing proposals to derive insights into how citizens view particular government decisions or public policies [16, 17, 23, 28].

As part of our future work, we plan to apply the proposed methodology to other polarized topics such as "legalization of abortion" and "right to bear arms". To this end, it will be necessary to collect and label a set of tweets related to these issues, which will make it possible to train issue-specific stance detection models. Finally, we plan to address usability aspects to integrate the proposed methods into a full-fledged software tool that can be used to critically explore polarized topics.

ACKNOWLEDGMENTS

This work was enabled by support provided by CONICET, Universidad Nacional del Sur (PGI-UNS 24/N051) and ANPCyT (PICT 2018-03627 and PICT 2019-03944). We thank anonymous reviewers for their constructive comments.

REFERENCES

[1] Lada A Adamic and Natalie Glance. 2005. The political blogosphere and the 2004 US election: divided they blog. In *Proceedings of the 3rd international workshop on Link discovery*. 36–43.

[2] Abeer ALDayel and Walid Magdy. 2021. Stance detection on social media: State of the art and trends. *Information Processing & Management* 58, 4 (2021), 102597.

[3] Mohammed N. Alenezi and Zainab M. Alqenaei. 2021. Machine Learning in Detecting COVID-19 Misinformation on Twitter. *Future Internet* 13, 10 (2021), 244. https://doi.org/10.3390/fi13100244

[4] Alexandra Balahur. 2013. Sentiment analysis in social media texts. In *Proceedings of the 4th workshop on computational approaches to subjectivity, sentiment and social media analysis*. 120–128.

[5] Carlos Iván Chesñevar, María Paula González, Ana Gabriela Maguitman, and Elsa Estevez. 2020. A first approach towards integrating computational argumentation in cognitive cities. In *ICEGOV 2020: 13th International Conference on Theory and Practice of Electronic Governance, Athens, Greece, 23-25 September, 2020*, Yannis Charalabidis, Maria Alexandra Cunha, and Demetrios Sarantis (Eds.). ACM, 25–32. https://doi.org/10.1145/3428502.3428506

[6] Jacob Cohen. 1960. A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement* 20, 1 (1960), 37–46. https://doi.org/10.1177/001316446002000104 arXiv:https://doi.org/10.1177/001316446002000104

[7] Michael Conover, Jacob Ratkiewicz, Matthew Francisco, Bruno Gonçalves, Filippo Menczer, and Alessandro Flammini. 2011. Political polarization on twitter. In *Proceedings of the international aaai conference on web and social media*, Vol. 5. 89–96.

[8] Kristijonas Cyras, Antonio Rago, Emanuele Albini, Pietro Baroni, and Francesca Toni. 2021. Argumentative XAI: A Survey. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, Zhi-Hua Zhou (Ed.). ijcai.org, 4392–4399. https://doi.org/10.24963/ijcai.2021/600

[9] Leon Derczynski, Torben Oskar Albert-Lindqvist, Marius Venø Bendsen, Nanna Inie, Viktor Due Pedersen, and Jens Egholm Pedersen. 2019. Misinformation on Twitter During the Danish National Election: A Case Study. In *Proceedings of the 2019 Truth and Trust Online Conference (TTO 2019), London, UK, October 4-5, 2019*, Maria Liakata and Andreas Vlachos (Eds.). https://truthandtrustonline.com/wp-content/uploads/2019/09/paper_16.pdf

[10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).

[11] Zulfadzli Drus and Haliyana Khalid. 2019. Sentiment analysis in social media and its application: Systematic literature review. *Procedia Computer Science* 161 (2019), 707–714.

[12] Shalmoli Ghosh, Prajwal Singhania, Siddharth Singh, Koustav Rudra, and Saptarshi Ghosh. 2019. Stance detection in web and social media: a comparative study. In *International Conference of the Cross-Language Evaluation Forum for European Languages*. Springer, 75–87.

[13] Kathrin Grosse, Carlos Iván Chesñevar, and Ana Gabriela Maguitman. 2012. An Argument-based Approach to Mining Opinions from Twitter. In *AT*. 408–422. http://ceur-ws.org/Vol-918/111110408.pdf

[14] Kathrin Grosse, Maria P Gonzalez, Carlos I Chesnevar, and Ana G Maguitman. 2015. Integrating argumentation and sentiment analysis for mining opinions from Twitter. *AI Communications* 28, 3 (2015), 387–401.

[15] Kadhim Hayawi, Sakib Shahriar, Mohamed Adel Serhani, Ikbal Taleb, and Sujith Samuel Mathew. 2022. ANTi-Vax: a novel Twitter dataset for COVID-19 vaccine misinformation detection. *Public health* 203 (2022), 23–30.

[16] Rocío B. Hubert, Elsa Estevez, Ana Maguitman, and Tomasz Janowski. 2018. Examining Government-citizen Interactions on Twitter Using Visual and Sentiment Analysis. In *Proceedings of the 19th Annual International Conference on Digital Government Research: Governance in the Data Age* (Delft, The Netherlands) *(dgo '18)*. ACM, New York, NY, USA, Article 55, 10 pages. https://doi.org/10.1145/3209281.3209356

[17] Rocío B Hubert, Elsa Estevez, Ana Maguitman, and Tomasz Janowski. 2020. Analyzing and Visualizing Government-Citizen Interactions on Twitter to Support Public Policy-making. *Digital Government: Research and Practice* 1, 2 (2020), 1–20. https://dl.acm.org/doi/abs/10.1145/3360001

[18] Clayton Hutto and Eric Gilbert. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the international AAAI conference on web and social media*, Vol. 8. 216–225.

[19] Zaid Jilani and Jeremy Adam Smith. 2019. What is the true cost of polarization in america? *Greater Good Magazine, March* 4 (2019).

[20] Dilek Küçük and Fazli Can. 2020. Stance detection: A survey. *ACM Computing Surveys (CSUR)* 53, 1 (2020), 1–37.

[21] Ema Kusen and Mark Strembeck. 2018. Politics, sentiments, and misinformation: An analysis of the Twitter discussion on the 2016 Austrian Presidential Elections. *Online Soc. Networks Media* 5 (2018), 37–50. https://doi.org/10.1016/j.osnem.2017.12.002

[22] Christian E. López, Malolan Vasu, and Caleb Gallemore. 2020. Understanding the perception of COVID-19 policies by mining a multilanguage Twitter dataset. *CoRR* abs/2003.10359 (2020). arXiv:2003.10359 https://arxiv.org/abs/2003.10359

[23] Asdrúbal López-Chau, David Valle-Cruz, and Rodrigo Sandoval-Almazán. 2020. Sentiment analysis of Twitter data through machine learning techniques. In *Software engineering in the era of cloud computing*. Springer, 185–209.

[24] Arjun Magge, Varad Pimpalkhute, Divya Rallapalli, David Siguenza, and Graciela Gonzalez-Hernandez. 2020. UPennHLP at WNUT-2020 Task 2 : Transformer models for classification of COVID19 posts on Twitter. In *Proceedings of the Sixth Workshop on Noisy User-generated Text, W-NUT@EMNLP 2020 Online, November 19, 2020*, Wei Xu, Alan Ritter, Tim Baldwin, and Afshin Rahimi (Eds.). Association for Computational Linguistics, 378–382. https://doi.org/10.18653/v1/2020.wnut-1.52

[25] Saif 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.

[26] Federico Neri, Carlo Aliprandi, Federico Capeci, Montserrat Cuadros, and Tomas By. 2012. Sentiment analysis on social media. In *2012 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. IEEE, 919–926.

[27] Orestis Papakyriakopoulos and Ellen Goodmann. 2022. The Impact of Twitter Labels on Misinformation Spread and User Engagement: Lessons from Trump's Election Tweets. In *WWW '22: The ACM Web Conference 2022, Virtual Event, Lyon, France, April 25 - 29, 2022*, Frédérique Laforest, Raphaël Troncy, Elena Simperl, Deepak Agarwal, Aristides Gionis, Ivan Herman, and Lionel Médini (Eds.). ACM, 2541–2551. https://doi.org/10.1145/3485447.3512126

[28] Rodrigo Sandoval-Almazan and David Valle-Cruz. 2020. Sentiment analysis of facebook users reacting to political campaign posts. *Digital Government: Research and Practice* 1, 2 (2020), 1–13.

[29] Parinaz Sobhani. 2017. *Stance detection and analysis in social media*. Ph.D. Dissertation. Universite d'Ottawa/University of Ottawa.

[30] Maxwell A. Weinzierl and Sanda M. Harabagiu. 2022. Identifying the Adoption or Rejection of Misinformation Targeting COVID-19 Vaccines in Twitter Discourse. In *WWW '22: The ACM Web Conference 2022, Virtual Event, Lyon, France, April 25 - 29, 2022*, Frédérique Laforest, Raphaël Troncy, Elena Simperl, Deepak Agarwal, Aristides Gionis, Ivan Herman, and Lionel Médini (Eds.). ACM, 3196–3205. https://doi.org/10.1145/3485447.3512039

[31] Sarita Yardi and Danah Boyd. 2010. Dynamic debates: An analysis of group polarization over time on twitter. *Bulletin of science, technology & society* 30, 5 (2010), 316–327.

[32] Lin Yue, Weitong Chen, Xue Li, Wanli Zuo, and Minghao Yin. 2019. A survey of sentiment analysis in social media. *Knowledge and Information Systems* 60, 2 (2019), 617–663.