

Applying the Ego Network Model to Cross-Target Stance Detection

Jack Tacchi^{*†§}, Parisa Jamadi Khiabani^{†§}, Arkaitz Zubiaga[†], Chiara Boldrini^{*}, Andrea Passarella^{*}

^{*} *Istituto di Informatica e Telematica, Consiglio Nazionale delle Ricerche, Pisa, Italy*

[†] *School of Electronic Engineering and Computer Science, Queen Mary University, London, United Kingdom*

[‡] *Scuola Normale Superiore, Pisa, Italy*

Abstract—Understanding human interactions and social structures is an incredibly important task, especially in such an interconnected world. One task that facilitates this is Stance Detection, which predicts the opinion or attitude of a text towards a target entity. Traditionally, this has often been done mainly via the use of text-based approaches, however, recent work has produced a model (CT-TN) that leverages information about a user’s social network to help predict their stance, outperforming certain cross-target text-based approaches. Unfortunately, the data required for such graph-based approaches is not always available. This paper proposes two novel tools for Stance Detection: the Ego Network Model (ENM) and the Signed Ego Network Model (SENM). These models are founded in anthropological and psychological studies and have been used within the context of social network analysis and related tasks (e.g., link prediction). Stance Detection predictions obtained using these features achieve a level of accuracy similar to the graph-based features used by CT-TN while requiring less and more easily obtainable data. In addition to this, the performances of the inner and outer circles of the ENM, representing stronger and weaker social ties, respectively are compared. Surprisingly, the outer circles, which contain more numerous but less intimate connections, are more useful for predicting stance.

Index Terms—stance detection, cross-target, ct-tn, ego network, signed relationships, social media

I. INTRODUCTION

Humans have always been social animals. Everyday we interact countless times with one another and these interactions form the basis of our modern societies. It has even been argued that our ability to maintain larger social groups was the primary reason that the volume of our brains increased significantly in size [1]. What’s more, since the advent of the internet, humans have also been able to interact with each other regardless of geographical location. Because humans are now more interconnected than ever, understanding social connections and how they contribute to the spreading of ideas and opinions has never been more important.

The exponential expansion of social networks has introduced fresh obstacles in the realm of information retrieval [2]. One research task that deals directly with opinions on social networks is Stance Detection (SD), which aims to predict the stance of a given text towards a target entity [3]. Of course, being able to monitor the opinions of individuals or even overall trends in larger communities and populations can be extremely impactful. Especially, given its application to politics, where it can be used to quickly understand how people feel towards a given topic or to predict how they will vote.

SD has often viewed interactions in isolation, predicting a user’s opinion towards a given entity purely based on what they have written (e.g. [4]). However, recent research has shown that considering an individual’s surrounding social network can greatly improve the accuracy of SD [5], highlighting the influence of social connections on opinions. Specifically, a model called CT-TN (Cross-Target Text-Net), which uses predictions from a more traditional text-based model, RoBERTa [6], together with multiple network features from the X (formerly Twitter) social media platform: likes (a list of users whose posts have been liked by the target user), followers (the users who follow the target user) and friends (the users who are followed by the target user). This was done by creating an embedding for each feature and passing it through a classification model to obtain a prediction for each feature. The final prediction was then generated based on a majority vote of all the features (see Figure 1). The CT-TN model outperformed other competitive models, such as CrossNET [7] and TGA-Net [8] in six different experimental conditions.

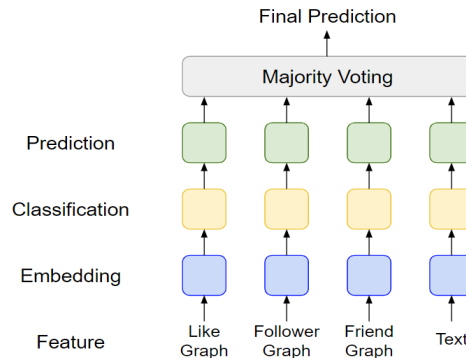


Fig. 1. Architecture of the CT-TN model.

However, the different network features that CT-TN requires are not always available. Indeed, the multiple different data sources required for CT-TN may be impossible or extremely costly to obtain in many situations. Therefore, it would be pertinent to investigate alternative approaches or features that are more parsimonious.

To this effect, in this work, we exploit models of users’ personal networks, grounded on well-established findings in anthropology. In addition to providing solid quantitative models of the structure of users’ social relationships, which can be obtained whenever users communicate publicly, thus minimising

[§]Co-first authors with equal contribution and importance.

restrictions from the aforementioned shortcomings as much as possible. Indeed, given the aforementioned importance of understanding how humans communicate, it is unsurprising that this topic had been researched long before the internet. For instance, anthropological and psychological studies have found that the number of relationships that an individual is able to maintain is remarkably consistent across members of the same species [9]. The maximum maintainable group size is correlated with the proportional size of the neocortex part of said species’ brain; strongly suggesting that a group’s size is innately limited by the cognitive ability of the species within it [1]. What’s more, these connections can be sorted by contact frequency into a series of concentric groups [10], with a smaller number of high-frequency connections at the centre and larger numbers of less-frequent connections towards the edges.

A representation of this structure, known as the Ego Network Model (ENM), places a target individual (the Ego, from which the model takes its name) at the centre and surrounds them by all of their connections (Alters) (see Figure 2). For humans, the expected sizes of the concentric groups range from 5 (support clique), to 15 (sympathy group), then around 45-50 (affinity group) and finally 150 (active network) [10]. An additional inner circle of 1-2 connections has also been found in many online contexts [11]. Interestingly, these groups increase in size by a factor of around 3 each time, and this scaling factor has also been found for certain non-human animals, such as other primates and even birds [12].

The ENM has been studied not only offline but has also been used to reveal many novel insights about communications in online contexts [13]. For instance, the ENM has been used to discover differences between the behaviours of various types of users, including journalists [14] and members of different online communities [15]. It has also been shown that using ENM information can significantly improve “classical” tasks in Online Social Networks (OSNs), such as link prediction [16]. Moreover, recent extensions of ENM have been proposed to label links in the Ego Network with positive or negative signs, depending on whether the corresponding relationship between Ego and Alter displays a predominantly positive or negative sentiment [17]. This extension, called the Signed ENM (SENM) has been used to characterise, in detail, the prevailing polarity of social relationships in OSNs, in contrast to the offline world [18]. Therefore, the ENM and SENM could reasonably be expected to provide significant information pertinent to the task of SD. Indeed, Ego Networks go beyond providing mere structural information but can contribute semantically to a better knowledge of the underlying mechanisms of a network and, therefore, provide pertinent contextual insights for SD. In addition, the base ENM requires only a list of Ego-Alter pairs and the frequencies of their interaction; data which can be easily obtained just by monitoring public interactions.

A. Contributions

Given the potential restrictions of data availability discussed above, it is important to have a toolbox of diverse approaches




Fig. 2. The Ego Network Model, with the expected number of Alters of each circle (for humans).

to ensure that researchers have a viable means of tackling as many situations as possible. Indeed, the authors of the CT-TN model specifically advocate for exploring further network features to enhance SD [5]. In response to this, the current work proposes two features from adjacent network research areas for use within the task of SD: the ENM and the SENM. Furthermore, this paper demonstrates that these features perform better than a text-only approach, RoBERTa, and similarly, although slightly worse, than the cutting-edge CT-TN model. Thus, demonstrating that these novel features are viable alternatives for SD when not all the required features for CT-TN can be obtained. Finally, a comparison of the performances between the inner and outer circles provides evidence that a user’s less intimate but more numerous outer connections have more of an impact on their stance.

II. BACKGROUND

A. Stance Detection

The focus of most SD studies is to classify target dependency in one of four main ways: Target-Specific, Multi-Related Targets, Target-Independent and Cross-Target [19]. For this paper, we primarily concentrate on Cross-Target Stance Detection (CTSD), which is when a model is trained using data for one target entity (source) and then tested on a different, although related, target entity (destination). For example, a model trained using texts containing opinions towards Joe Biden could be used to predict the stance of texts concerning another politician, such as Donald Trump or Bernie Sanders.

Approaches to CTSD, as well as to SD more generally, differ based on the text’s context and the particular relationship being discussed. On social media platforms such as X, there’s often a focus on discerning the author’s stance (supportive, opposing, or neutral) towards a specific proposition or target [20]. Recent advances in SD encompass a range of linguistic features, such as word or character n-grams, dependency parse trees, and lexicons [21]. Moreover, there has been a shift towards end-to-end neural network methods that independently learn topics and opinions, integrating them via mechanisms such as memory networks or bidirectional conditional Long Short-Term Memory models (LSTMs) [22].

Past studies have primarily relied on the text of a post to gauge its stance, neglecting the valuable insights that other features within social media platforms could offer. However,

the performance of the aforementioned CT-TN model demonstrates the importance of considering structural features of the surrounding social network. Thus, knowledge about a user’s connections can reveal important insights about the user themselves and, therefore, about the texts they author.

B. The Ego Network Model

As mentioned in Section I, the ENM views a social network from the point of view of an individual Ego and organises their Alters around them based on their contact frequency, while the SENM extends this by adding the addition of signed relationships. Specifically, these signed relationships are obtained by analysing the sentiment of the interactions between each Ego and Alter, thus obtaining a list of sentiments for each relationship, which can then be used to infer an overall sentiment for the relationship as a whole. It has been observed that many different types of relationships start to have negative effects on the people involved once the ratio of negative interactions generated by that relationship passes a certain ratio: around 17% [23]. Therefore, this threshold is used to infer the sentiments of relationships based on their interactions.

While the Signed Ego Network provides an additional layer of information, it does require text for each interaction in a dataset. By contrast, the unsigned Ego Network requires only the frequency of interactions between each pair of users, without the need for any text.

III. METHODOLOGY

A. Performing Cross-Target Stance Detection

1) *Feature Embeddings*: In order to compute SD predictions, the data first need to be transformed into representations that are readable by a prediction model. For this, node2Vec [24] was applied to each of the previously established graph-based features (likes, followers, friends) as well as the novel Ego Networks and Signed Ego Networks, which, although they are converted into the same vector-space representation, can be better thought of as proxy measures of the way humans function socially. Node2vec is an unsupervised Deep Learning algorithm that uses a flexible, biased, random walk procedure to explore networks. The visited nodes can then be transformed into a vector space representation using a variety of methods, such as skip-grams or a continuous bag-of-words [24]. This is similar to how the word2vec algorithm [25] treats words (nodes) and sentences (walks).

In addition to the graph-based features, text-based predictions were also used. These were generated using RoBERTa [25], an incredibly well-performing pretrained model that is used for many different natural language processing tasks. RoBERTa maps every token in a sentence to a vector representation in a continuous space.

As mentioned in Section I, the CT-TN model takes the predictions of each of these features, RoBERTa, likes, followers and friends, and obtains a final prediction using majority voting, where each feature’s prediction acts as a vote for either “FAVOR” or “AGAINST”. This allows for a thorough analysis

of both textual and social network information, providing valuable insights for CTSD.

Aside from the aforementioned features, this paper also investigates two novel graph-based features: Ego Networks and Signed Ego Networks. These are also converted to a vector space representation using node2vec. Further details on how the Ego Networks are obtained are explained in Section III-B.

2) *Model Hyperparameters*: Each of the features was used to train a neural network model with two hidden layers for classification task. For the text-based embedding, this was done using RoBERTa, a batch size of 128, a dropout of 0.2, a learning rate of $3e-5$ (AdamW), and 40 epochs. For the graph-based embeddings, this was done using node2vec, with the same batch size and dropout, a learning rate of $1e-2$ (SGD), and 100 epochs.

3) *Experimental Settings*: The CT-TN model and the individual RoBERTa feature predictions were used as baselines against which to test the two Ego Network features. These were all prepared using few-shot cross-target training, whereby the training data consisted of roughly 1,000 source target data points with 4 injections of destination target texts, increasing in size by increments of 100, from 100-shot to 400-shot (inclusive). For example, the Biden-Trump predictions were obtained by training on around 1,000 Biden texts with 100 Trump texts for the 100-shot condition, with 200 Trump tweets for the 200-shot condition, and so on. The stance predictions were then tested using between 500 and 800 data points (depending on the amount of remaining unseen data) that were solely related to the destination target (i.e. only using Trump-related texts for the aforementioned example). We conducted these experiments with five different random seeds: 24, 524, 1024, 1524, and 2024. Finally, we averaged the results from these five seeds for each shot size.

B. Computing Ego Networks

1) *Computing Ego Networks*: An Ego Network can be obtained for each active¹ user in the data by calculating the interaction frequencies of each of their relationships and then applying a clustering algorithm to them. One of the most popular methods for this is MeanShift [26], an unsupervised algorithm that automatically finds the most appropriate number of circles for an Ego [13], and so that is the algorithm that was employed for this work.

Additionally, the unsigned Ego Network feature was also separated into inner circles (1 and 2) and outer circles (3+), to better understand the importance of the different levels of the ENM for SD.

2) *Generating Signed Connections*: Once the unsigned Ego Networks have been obtained, it is relatively simple to generate the signed version. The interactions were grouped for each Ego-Alter pair and then a model, called Valence Aware Dictionary and sEntiment Reasoner (VADER) [27], was used to obtain a sentiment for each interaction. VADER is a very

¹An active user is defined as a user, with a timeline of at least 6 months, that posts at least once every three days for half of the months that they are included in the data, based on previous research [13]

competitive sentiment analysis model that is specifically designed for performing sentiment analysis on English-language tweets. It is also the model used in the original SENM research paper [17].

Once sentiments had been obtained for the individual interactions, the psychology-based threshold of 17% [23], mentioned in Section II-B, was used to obtain a sentiment label for each relationship, resulting in signed Ego Networks.

C. Data

The data used in this study come from a well-established and publicly available set: P-Stance [28]. This dataset contains 21,574 English-language tweets collected during the 2020 U.S. presidential election. These tweets were specifically collected to be used for SD and each one is associated with one of 3 targets, Joe Biden, Donald Trump or Bernie Sanders, and a corresponding stance label, either “AGAINST” or “FAVOR”. Hashtags, such as “#BidenForPresident” and “#NeverBernie” were used to both search for the tweets and to determine their target and stance. While the original dataset only included the text and stance of each tweet, the authors of the P-Stance dataset provided 9,307 tweet IDs upon request, allowing further data to be collected for each tweet, including information about the authoring users. In addition to providing the information required for computing the users’ Ego Networks (see Section III-B), this also made it possible to obtain, for each user, the remaining features required by CT-TN: likes, followers and friends.

IV. RESULTS

The performances of the CT-TN model, RoBERTa, and the two Ego Network features can be seen in Figure 3. Overall, they all perform very well, with most reaching a macro F1 score of above 0.7 before 400-shot for all target pairs, with CT-TN sometimes even going above 0.8. However, RoBERTa does not perform quite as well as the others, and only achieves macro F1 scores of around mid-0.6 for half of the target pairs.

Surprisingly, the signed and unsigned Ego Networks’ F1 scores are very close, being within 0.01 of each other for 5 of the 6 target pairs (at 400-shot), and within 0.02 for the sixth pair (Sanders-Trump). This suggests that the additional information of signed connections does not provide a significant amount of information for the task of CTSD. Rather, it appears that the people we interact with regularly have an impact on our stances regardless of whether we have a negative or positive relationship with them.

Next, observing the outer circles, one can see that they perform similarly to, and often even outperform, the full Ego Network. However, they are less consistent, as displayed by the Biden-Trump and Sanders-Trump target pairs. By comparison, the inner circles perform slightly worse overall, with performances closer to those of RoBERTa. Since the outer circles contain weaker social relationships, it seems that weaker, but more numerous, ties are more informative than stronger, but less numerous, ones when it comes to stance prediction. This is rather surprising given that previous

research on a similar network-based task, link prediction, found that the more intimate inner circles are better predictors of where new relations will form [16]. Thus, there seems to be a disconnect between how we form new connections and how we are influenced by them. Indeed, paired with the fact that the signed and unsigned ENMs performed similarly, it appears that the existence of a social connection may influence an individual regardless of any qualitative aspects, such as closeness or polarity.

The Ego Networks appear to perform slightly worse than the CT-TN model. However, as they only require interaction data, they could be used as a viable alternative whenever specific network features are not provided or obtainable for a given dataset. Moreover, as the signed and unsigned Ego Networks achieved similar performances, one could focus on employing the unsigned version, which would require even less data: only the frequencies of interactions, without the need for their texts.

V. CONCLUSION

This study has highlighted potential limitations with previous approaches to CTSD (and SD more generally) due to the availability, or rather the unavailability, of data. This is especially relevant given recent restrictions to data accessibility on X, most notably the discontinuation of the Academic API, which has been one of the largest and most popular social network data sources for academic use for over a decade [11], [13], [17], [18].

Addressing this limitation, this paper has proposed two novel graph-based features that can be used for SD: the ENM and the SENM. The latter requires only text-based interactions between users and the former only requires interaction frequencies, meaning that the content of the interactions can remain hidden. These features perform consistently well for CTSD (achieving macro F1 scores of at least 0.71, after 400 shots, for all 6 target pairs tested in this paper). Their performances are only slightly worse than a previously established and cutting-edge model, CT-TN. Ego Networks present viable alternatives that could be used when not all the features required for CT-TN (or similar models) are obtainable.

Finally, by observing the different performances between the inner and outer circles of the ENM, it appears that, while the inclusion of both leads to a more consistent performance across different target pairs, the outer circles on their own can often perform just as well and sometimes better. By contrast, the inner circles do not perform as well, suggesting that the greater number of less intimate connections in the outer circles are more important for predicting a user’s stance.

REFERENCES

- [1] R. I. M. Dunbar, “The social brain hypothesis,” *Evolutionary Anthropology: Issues, News, and Reviews: Issues, News, and Reviews*, vol. 6, no. 5, pp. 178–190, 1998.
- [2] P. J. Khiabani, M. E. Basiri, and H. Rastegari, “An improved evidence-based aggregation method for sentiment analysis,” *Journal of Information Science*, vol. 46, no. 3, pp. 340–360, 2020.
- [3] D. Biber and E. Finegan, “Adverbial stance types in english,” *Discourse processes*, vol. 11, no. 1, pp. 1–34, 1988.

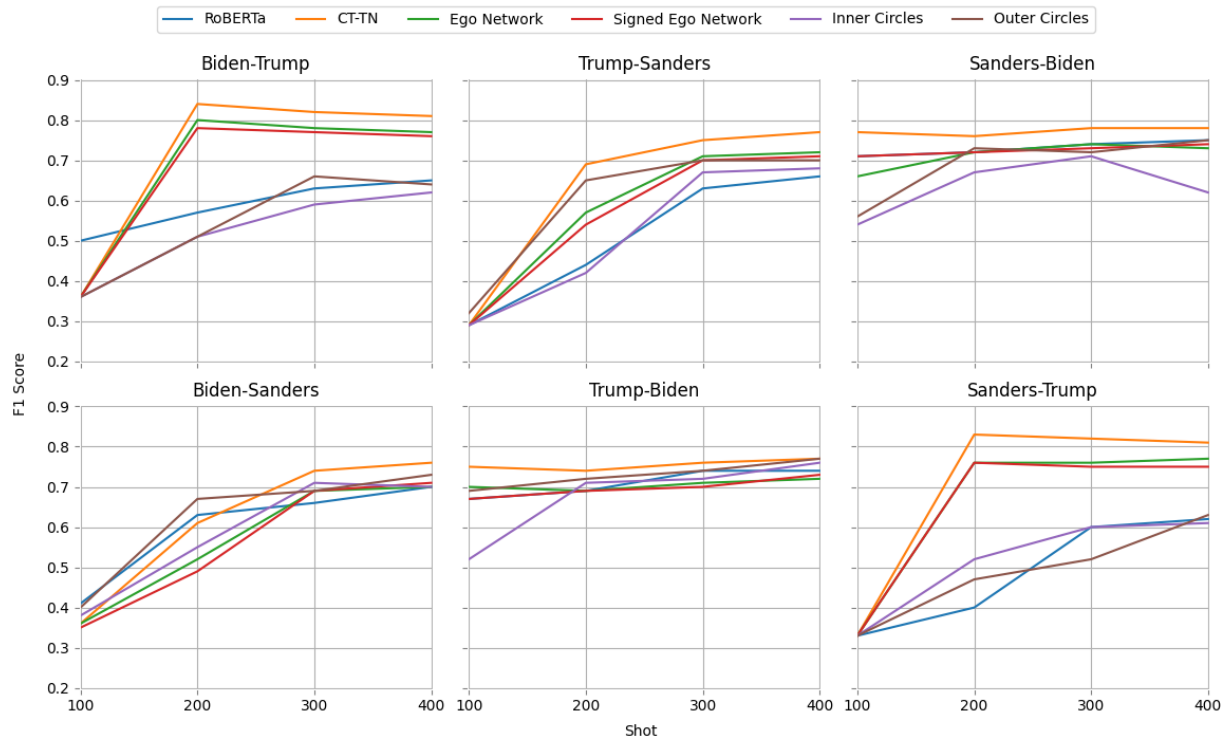


Fig. 3. Graphs displaying the performances of the individual text-based RoBERTa, the CT-TN model and the signed and unsigned Ego Networks features, as well as the unsigned inner and outer circle features. The shot number is displayed along the X axis and the averaged macro F1 scores along the Y axis.

- [4] P. Wei and W. Mao, "Modeling transferable topics for cross-target stance detection," in *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2019, pp. 1173–1176.
- [5] P. J. Khiabani and A. Zubiaga, "Few-shot learning for cross-target stance detection by aggregating multimodal embeddings," *IEEE Transactions on Computational Social Systems*, 2023.
- [6] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "Roberta: A robustly optimized bert pretraining approach," *arXiv preprint arXiv:1907.11692*, 2019.
- [7] C. Xu, C. Paris, S. Nepal, and R. Sparks, "Cross-target stance classification with self-attention networks," *arXiv preprint arXiv:1805.06593*, 2018.
- [8] E. Allaway and K. McKeown, "Zero-shot stance detection: A dataset and model using generalized topic representations," *arXiv preprint arXiv:2010.03640*, 2020.
- [9] R. I. M. Dunbar, "Neocortex size as a constraint on group size in primates," *Journal of human evolution*, vol. 22, no. 6, pp. 469–493, 1992.
- [10] R. I. M. Dunbar and M. Spoors, "Social networks, support cliques, and kinship," *Human nature*, vol. 6, pp. 273–290, 1995.
- [11] V. Arnaboldi, M. Conti, A. Passarella, and F. Pezzoni, "Ego networks in twitter: an experimental analysis," in *2013 Proceedings IEEE INFOCOM*. IEEE, 2013, pp. 3459–3464.
- [12] R. I. M. Dunbar *et al.*, *Mind the gap: or why humans aren't just great apes*. Oxford University Press, 2014.
- [13] C. Boldrini, M. Toprak, M. Conti, and A. Passarella, "Twitter and the press: an ego-centred analysis," in *Companion Proceedings of the The Web Conference 2018*, 2018, pp. 1471–1478.
- [14] M. Toprak, C. Boldrini, A. Passarella, and M. Conti, "Structural models of human social interactions in online smart communities: the case of region-based journalists on twitter," *Online Social Networks and Media*, vol. 30, 2021.
- [15] J. Tacchi, C. Boldrini, A. Passarella, and M. Conti, "Cultural differences in signed ego networks on twitter: An investigatory analysis," in *Companion Proceedings of the ACM Web Conference 2023*, 2023, pp. 1039–1049.
- [16] M. Toprak, C. Boldrini, A. Passarella, and M. Conti, "Harnessing the power of ego network layers for link prediction in online social networks," *IEEE Trans. Comput. Soc. Syst.*, vol. 10, no. 1, pp. 48–60, 2023.
- [17] J. Tacchi, C. Boldrini, A. Passarella, and M. Conti, "Signed ego network model and its application to twitter," *IEEE BigData 2022*, 2022.
- [18] —, "On the joint effect of culture and discussion topics on x (twitter) signed ego networks," *arXiv preprint arXiv:2402.18235*, 2024.
- [19] N. Alturayeif, H. Luqman, and M. Ahmed, "A systematic review of machine learning techniques for stance detection and its applications," *Neural Computing and Applications*, vol. 35, no. 7, pp. 5113–5144, 2023.
- [20] S. Mohammad, S. Kiritchenko, P. Sobhani, X. Zhu, and C. Cherry, "Semeval-2016 task 6: Detecting stance in tweets," in *Proceedings of the 10th international workshop on semantic evaluation (SemEval-2016)*, 2016, pp. 31–41.
- [21] S. Sun, C. Luo, and J. Chen, "A review of natural language processing techniques for opinion mining systems," *Information fusion*, vol. 36, pp. 10–25, 2017.
- [22] I. Augenstein, T. Rocktäschel, A. Vlachos, and K. Bontcheva, "Stance detection with bidirectional conditional encoding," *arXiv preprint arXiv:1606.05464*, 2016.
- [23] J. Gottman, J. M. Gottman, and N. Silver, *Why marriages succeed or fail: And how you can make yours last*. Simon and Schuster, 1995.
- [24] A. Grover and J. Leskovec, "node2vec: Scalable feature learning for networks," in *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 2016, pp. 855–864.
- [25] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *arXiv preprint arXiv:1301.3781*, 2013.
- [26] K. Fukunaga and L. Hostetler, "The estimation of the gradient of a density function, with applications in pattern recognition," *IEEE Transactions on information theory*, vol. 21, no. 1, pp. 32–40, 1975.
- [27] C. Hutto and E. Gilbert, "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, no. 1, 2014, pp. 216–225.
- [28] Y. Li, T. Sosea, A. Sawant, A. J. Nair, D. Inkpen, and C. Caragea, "P-stance: A large dataset for stance detection in political domain," in *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 2021, pp. 2355–2365.