Assessing the Effects of Social Familiarity and Stance Similarity in Interaction Dynamics

Kuntal Dey[1(✉)], Ritvik Shrivastava[2], Saroj Kaushik[3], and Vaibhav Mathur[2]

[1] IBM Research, New Delhi, India
kuntadey@in.ibm.com
[2] Netaji Subhas Institute of Technology, New Delhi, India
{ritviks.it,vaibhavm.it}@nsit.net.in
[3] Indian Institute of Technology, Delhi, India
saroj@cse.iitd.ac.in

Abstract. Homophily, the phenomenon of similar people getting connected to and being socially familiar with each other, is well-known on online social networks. Detection of user stance towards given topics, on online social networks, specifically Twitter, has emerged as a mainstream research topic. The current work provides insights into the impact of topic-specific stance similarity and social familiarity on social interaction dynamics. This is a novel and yet fundamental problem in social networks research, that has so far remained unexplored in the literature. Specifically, we address two key aspects. One, we investigate whether the smoothness (politeness) level of conversations between user pairs, relate with overall stance similarity (spanning across topics). Two, we examine the impact on interaction smoothness (politeness) with respect to social familiarity and topical stance-similarity. We propose a novel approach based on word embedding, to compare across users and across topics. We analyze the relationship between topical stance similarity, social familiarity and interaction politeness of users, with respect to specific familiarities between user pairs as well as social communities.

1 Introduction

Similarity and familiarity are two core properties of individuals participating in online social networks. Homophily [19] is based on the hypothesis: *similarity breeds connection* (familiarity). As observed by the seminal work, "homophily limits peoples social worlds in a way that has powerful implications for the information they receive, the attitudes they form, and the interactions they experience". Clearly, it is a fundamental property of social networks.

Of late, stance detection with respect to given topics, on online social networks, specifically Twitter, has emerged as a mainstream research work, motivated by the SemEval 2016 challenge [22,23]. Stance detection addresses the problem of understanding the user's stand - favoring, neutral or against - a given topic. Several works, such as [6,34,37,38] to name a few, have emerged addressing user stance detection toward given topics.

C. Cherifi et al. (eds.), *Complex Networks & Their Applications VI*,
Studies in Computational Intelligence 689, https://doi.org/10.1007/978-3-319-72150-7_68

In Twitter, while topic detection has remained a general area of research till date, hashtags have clearly been treated as one of the most popular markers of topics [9]. If hashtags are treated as topics, applying the stance detection models that are present in the literature, would yield the stance of a given user towards the hashtags discussed in a given tweet. This, along with the recently emerging research problem of user stance detection towards Twitter topics, creates a novel problem area in turn: *What do the analysis of user interaction dynamics indicate (subject to social familiarity and stance similarity metrics), when stances are determined as user's stand - positive, neutral or negative - with respect to a given hashtag-topic?* This question is the first-of-its-kind in research. Further, it creates the possibility of obtaining a finer-grained understanding of homophily (where similarity is determined as stance similarity with respect to specific topics), as compared to the state-of-the-art coarse grained insights of today (where similarity is determined as sentiment or explicit interest or activity based similarity).

Specifically, to establish a baseline for this novel research paradigm, we focus the current scope of our work towards two fundamental questions. First: *Does the politeness level of conversations between user pairs vary with overall stance similarity (spanning across topics)?* To this, we propose a latent concept space embedding based framework, for obtaining stance similarity of user pairs with respect to given topics (hashtags). Second: *What is the impact on interaction smoothness (politeness) with respect to social familiarity and topical stance-similarity?* To this, we derive community-level user similarity and interaction friction level (politeness), using our framework. Please note that, while the first question above pertains to user similarity, the second question is central to the concept of homophily, but is framed to analyze in a more fine-grained manner compared to how the traditional literature addresses homophily.

The key contributions of the current paper are as follows.

- We propose a first-of-its-kind exploratory study on social networks, wherein we explore the characteristics of user similarity as well as homophily, with respect to the stances of users towards given topics.
- We utilize a high-dimensional latent concept space similarity, and thereby create a framework for measuring stance similarities of user pairs, across a given set of topics.
- Using the platform thus established, we explore topic-specific stance similarity and familiarity characteristics of social network users, from the angles of interaction friction (politeness) at individual communication levels, as well as at community levels.
- We observe the relationship between topical stance similarity, social familiarity and interaction politeness of users, both from the standpoint of specific familiarities between user pairs as well as from the level of structural social communities.

2 Related Work

We observe that the problem we are addressing is novel, and there is no work in the literature that we can refer to or compare our work with. In this section, we introduce background work to the reader, that address some of the individual components constituted as part of our work.

Homophily [19], a form of assortative social mixing of like-minded individuals, is based on the hypothesis: *similarity breeds connection* (familiarity). Several research works have investigated the impact of homophily in social networks. De Choudhury *et al.* [11] observe that homophily impacts the communication behavior of individuals, and that in turn affects the information propagation mechanisms on the social networks as a whole. Halberstam and Knight [15] show a strong presence of homophily among politically engaged users on Twitter, and observe that the rate of information diffusion, as well as exposure of user to information content, is significantly higher in larger homophilous groups. Other works, such as [1,33], also address homophily for modeling social information diffusion dynamics.

The literature for topic identification on Twitter has followed three different approaches. In one approach, the hashtags contained as part of tweet messages, are treated as topics. Works, such as [9], use this approach to associate topics with tweets. In another approach, a burst of keywords in a short span of time are identified, and each bursting keyword is treated as a topic. Works, such as [7,8,18], use this approach. And in a third approach, the latent semantic concepts of given tweets - often identified with sophisticated text-to-topic assignment techniques such as Latent Diriclet Allocation (LDA) [4] - are treated as topics, and the tweets that address these concept spaces are said to belong to these topics. Works, such as [17], follow this approach.

Literature has recently observed *semantic homophily*, a similarity measure for user pairs based upon semantic features of communication content [30], shape up communications on online social networks. In our work, we use semantic properties of the user-generated content to identify the stance of the users towards given topics (hashtags). The emergence of stance detection methodologies such as [6,13,34,37,38], following the SemEval challenge of topic-specific stance detection of users [22,23], motivates us to obtain insights into stance-specific homophily, wherein the derivation of stance uses tweet content semantic features.

Bengio *et al.* [2] introduced the concept of word embedding. Mikolov *et al.* [20] proposed the Word2Vec model, that extended over the core concept of word embedding, and has gained significant popularity due to its inherent capability of performing unsupervised training of the embeddings. Word embedding and its different derivations have become integral parts of a wide range of solutions, such as [3,21,35] to name a few. Amongst other related approaches, GloVe [28] based embedding has also found research traction.

Politeness (and abuse) detection from user-generated online content has been a topic of research interests. Prior works, such as [26,27], have addressed the problem. Mizil *et al.* [10] had recently proposed a lexical and syntactic feature based classifier for politeness detection, and observed that highly reputed users tend to be less polite on social platforms. Schmidt and Wiegand [31] provide a survey of the politeness and abusive language literature.

3 Our Approach

3.1 The Semantic Similarity Framework Using Content Embedding

We create a framework that will allow us to measure the similarity of user pairs with respect to given topics. Akin to the philosophy of much of the literature such as [9], we use hashtags as topics, with the argument that the hashtags used to the users are explicit intention of the key message of the tweet.

We observe from our dataset that, while a lot of users have common hashtags, a signifcant fraction of the connected users do not always use the same set of hashtags, even while tweeting on related topics. Not considering interactions where one user uses a hashtag that the other has never used, is lossy in nature. In order to avoid this loss of information, we follow a soft-cosine similarity based approach (described below), which, to the best of our knowledge, has not been used in the literature for user pair stance similarity measurement.

Overall stance of user with respect to each hashtag
The first step is to determine the overall stance of a given user, with respect to each given hashtag. We isolate all the tweets t_{ip} that a given user u_i posts. We detect stance s_{ikp} of user u_i using an existing methodology [13], for each tweet, with respect to each hashtag h_k (topic) present in the tweet t_{ip}. We assign $+1$ to all the *favor* stances, 0 to all the *neutral* stances and -1 to all the *against* stances. The *overall stance* S_{ik} of user u_i towards each given topic (hashtag) h_k is obtained by summing up her stance towards the topic, by examining all the individual posts t_{ip} she makes for the topic. This is given as:

$$S_{ik} = \sum_{p} s_{ikp} \tag{1}$$

Semantic standpoint of user with respect to each hashtag
The next step is to determine the semantic standpoint of each given user, with respect to each hashtag. Since, each user will potentially use multiple hashtags that other users she converses with will not use, this step will create a platform for us to subsequently leverage in finding user pair similarity over non-matching hashtags.

Let us explain with an intuitive example. Say, user A talks about #TENNIS and #BASKETBALL, user B talks about #GOLF and user C talks about #DONALDTRUMP. Here, all the users are using different hashtags, and hence, it is not feasible to directly measure user pair similarities between any of the pairs, $\{A, B\}$, $\{A, C\}$ or $\{B, C\}$. And yet, given that all of tennis, basketball and golf are sports, while Donald Trump is likely to be a politics topic, the likelihood of similarity of users A and B (those who address the sports topics) ought to be higher compared to the similarity of A and C, as well as the similarity of B and C.

In order to enable measuring similarities for such cases, one traditional approach would involve using Ontologies (such as the Wikipedia ontology) [29], and subsequently comparing the different relationships of the nodes in the ontologies as well as the relationships of different ancestor nodes. Other feature-based approaches for measuring semantic similarity of texts also exist, such as Dey *et al.* [12]. We use the concept of word embedding for our problem, as that is known to better capture the high dimensionality of underlying latent concepts of text, compared to all the other existing approaches, across different applications [24,39]. Specifically, we use the pre-trained word embedding vector from Google News[1].

Let a given tweet t_{ip} comprise of the words $\{w_{ip}\}$, along with the hashtag h_k. The set of words used by the user, across all the posts, that also contain this hashtag, is given by concatenating the content of all the tweets (except, we leave the hashtag out as the embedding of the hashtag will be expressed as a function of embedding of these associated words), akin to the approach of [14]:

$$\{W_i\} = \biguplus_{(p|h_k \in \{w_{ip}\})} \{w_{ip} - h_k\} \tag{2}$$

The embedding v_{wh_k} of each word $w \in W_i$ is found from the provided external word embedding resource. Subsequently, we compute an embedding V_h for each hashtag as a whole, using the embedding of the words that appear in the tweets containing the hashtag. This is obtained as the average embeddings of all the words that appear across all the posts containing the hashtag.

$$V_{h_k} = \frac{\sum\limits_{w \in W_i} (v_{wh_k})}{|\{W_i\}|} \tag{3}$$

In the above, $|\{W_i\}|$ denotes the total number of words used in the hashtag content across the posts (with repeated words retained). The computation is carried out for all hashtags $h_k \in H$, which yields the (word-averaged) embedding of each hashtag in the given set of tweets. Note that, words that are repeated within as well as across posts, are considered as many times as they appear; that is, the repeating behavior is retained, as this inherently assigns the necessary weight that the embedding merits with respect to how highly each word is used with respect to that hashtag.

Stance similarity computation for user pairs

To compute the stance similarity of user pairs, the semantic similarity of a each pair of hashtags that a user pair posts, is first computed as Euclidean distance between the embedding vectors of the pair of hashtags. This is combined with

[1] https://drive.google.com/file/d/0B7XkCwpI5KDYNlNUTTlSS21pQmM/edit?usp=sharing.

the stance of each user to the hashtag she posts, in form of *soft cosine similarity* [32] between user u_i and u_j, to derive the overall stance similarity of users across topics. Note that, while a traditional cosine similarity considers the vector space to comprise of completely independent features, a soft cosine similarity considers the similarity of features in the vector space model. In our setting, the embedding provides a vector space model representation of the hashtags, from which, the Euclidean distance based vector distance (similarity) is computed.

Say, user u_i posts hashtags $\{h_q\}$, and user u_j posts hashtags $\{h_r\}$. The stance $\{S_q\}$ of user u_i to $\{h_q\}$ and $\{S_r\}$ of user u_j to $\{h_r\}$ are known by the earlier step of stance computation. Let u_i post M hashtags and u_j post N hashtags overall across all their posts. Let e_{qr} be the Euclidean distance of embeddings of hashtag h_q and h_r. Then, the semantic similarity between u_i and u_j are given by applying the soft cosine computation methodology, as:

$$soft\_cosine(u_i, u_j) = \frac{\sum_{q,r=1}^{M,N} e_{qr}.S_q.S_r}{\sqrt{\sum_{q,r=1}^{M,N} e_{qr}.S_q.S_q}\sqrt{\sum_{q,r=1}^{M,N} e_{qr}.S_r.S_r}} \tag{4}$$

Intuitively, the above is simply a normalized version of:

{(sum over all hashtags h_q that user u_i posts to) (sum over all hashtags that h_r that user u_j posts to) {(Euclidean distance of embedding of h_q and h_r) $\times$ (stance of u_i towards h_q) $\times$ (stance of (u_j towards h_r)}}

This gives similarity of user pairs, with respect to their stances over all the topics (hashtags) they participate in. The numerical values lie between -1 to $+1$.

3.2 Pairwise Politeness vs. Stance Similarity

We now attempt to answer one of the questions we had posed early on: *Does the politeness level of conversations between user pairs vary in with overall stance similarity (spanning across topics)?*

In order to find politeness of a given tweet, we use a well-accepted existing solution, proposed by Mizil *et al.* [10]. Their approach is based on identifying lexical and syntactic features from social content generated by users, and performing classification using two different approaches: (a) a bag-of-words classifier and (b) a linguistically informed classifier that uses 20 different politeness strategies, such as GRATITUDE, GREETING, APOLOGIZING, DIRECT START, NEGATIVE LEXICON *etc.* They perform Support Vector Machine (SVM) based machine learning, and predict the class probability labels as *politeness scores* that lie between 0 and 1.

In order to measure the politeness of interaction between user pair (u_i, u_j), since the Twitter data under consideration does not have any explicit conversation model, we take a collection of all the tweets where any one of the two users *mentions* the other. We find the politeness score for each tweet in this collection, and average out to find the age mutual politeness between the user pair. The final output of this step is a tuple for each user pair, each consisting of <user pair stance similarity, user pair mutual politeness>.

3.3 Politeness vs. Stance Similarity Given Social Familiarity

We now revisit the other question we had posed: *What is the impact on interaction smoothness (politeness) with respect to social familiarity and topical stance-similarity?*

We design our methodologies to answer this in two different manners. One, how does social familiarity (any or both of the *following* and *followed-by* relationships in Twitter) correlate with stance similarity of users with respect to topics? Two, does participation in social communities have any further impact?

In order to answer this, we divide the politeness of user pair interactions into different segments, and correlate with different levels of user-pair similarities. We further study the user pairs that are connected to (familiar with) each other, versus not.

In order to understand whether participation in social communities also has an impact on user pair interactions (subject to stance similarity), we perform community discovery using modularity maximization techniques [25]. We subsequently compute stance similarity and interaction politeness characteristics, and these for communications that happen between familiar users (members) within a given community versus familiar users that belong to different communities.

4 Experiments

Dataset Description

We conduct experiments the Twitter dataset[2] made available by Yang and Leskovec [36], that contains around 20–30% of the total number of tweets made during the period of collection. We perform experiments using all the tweets made between June 13th and 30th, 2009. We obtain all the social relationships of the users active during this period, from an online resource[3] made available by Kwak *et al.* [16]. We retain only that tweets that have at least one hashtag, since we shall aim to find stances treating hashtags as topics, and discard the rest of the tweets. We also discard retweets and quoted tweets (wherein someone quotes someone else's tweet, possibly with a few additional words within what the 140-character limit permits), as they do not add value in the context of our experiments. The statistics of the dataset are presented in Table 1.

---

[2] https://snap.stanford.edu/data/twitter7.html.

[3] http://an.kaist.ac.kr/traces/WWW2010.html.

Table 1. Description of Twitter data used for our experiments.

Total num. of tweets	Num. tweets retained	Num. users retained	Num. edges retained	Avg. num. tweets/user	Avg. num. connections
18,572,084	5,785,344	117,701	4,973,218	11.38	42.25

Tools and Resources
We use several well-known tools and resources for our work. For word embeddings, we use the pre-trained Twitter-specific version of GloVe word embedding [28]. In order to find modularity maximization based communities, we make use of BGLL [5]. For finding the politeness of a given tweet, we use Stanford politeness API[4], which gives a politeness value between 0 and 1. For stance, we use an existing algorithm [37]. We use the NLTK stopword list[5] for removing stopwords.

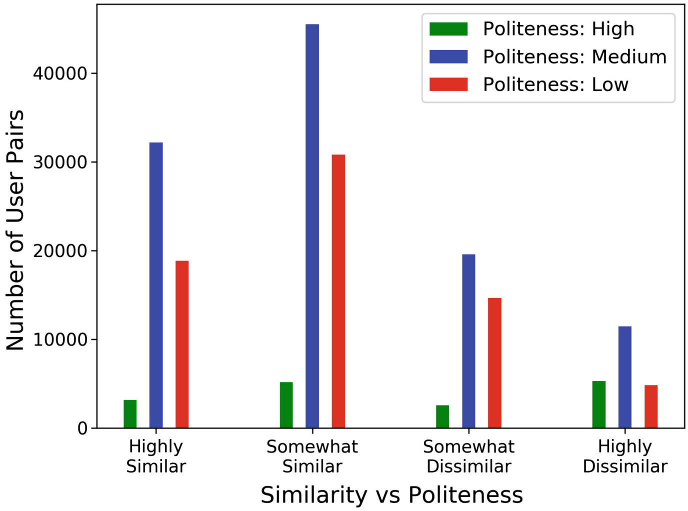

Fig. 1. Distribution of stance similarity and interaction politeness of user pairs

Collating Pairwise Politeness, Stance Similarity and Social Familiarity
Figure 1 captures the degrees of interaction politeness values for different levels of stance similarities, without considering familiarity. Table 2 presents the distribution of stance similarity and pairwise politeness, taking the presence and absence of familiarity into account. We make two interesting observations here.

---
4 https://github.com/sudhof/politeness.
5 https://github.com/stanfordnlp/CoreNLP/blob/master/data/edu/stanford/nlp/patterns/surface/stopwords.txt.

1. We observe that, highly dissimilar people are more polite to each other, irrespective of familiarity. Inspection of a random set of manual tweets suggest politeness arising from more formal messages (as agreed between two human annotators, when these random samples were presented for human validation post-experiments).
2. We also observe that, similar people tend to be more polite when they are familiar, which also in philosophy similar to the observations made by [10].

Table 3 present in order to give the reader an intuition of the data (such as, the hashtags the individuals use and their observed familiarity) and the results that our framework obtains (stance similarity and interaction politeness).

Table 2. Distribution of similarity, familiarity and politeness. "Highly similar" contain similarity values higher than 0.5, "somewhat similar" contain similarity values between 0 and 0.5, "somewhat dissimilar" contain similarity values between −0.5 and 0, and "highly dissimilar" contain similarity values smaller than −0.5.

Familiarity	Similarity	Politeness		
		High (> 0.67)	Medium (> 0.33 and < 0.67)	Low (< 0.33)
Yes	Highly similar	7.90%	57.41%	34.69%
	Somewhat similar	6.71%	59.91%	33.38%
	Somewhat dissimilar	6.30%	57.74%	35.96%
	Highly dissimilar	24.69%	58.26%	17.05%
No	Highly similar	5.71%	59.52%	34.77%
	Somewhat similar	6.33%	55.49%	38.18%
	Somewhat dissimilar	7.02%	52.83%	40.15%
	Highly dissimilar	24.52%	52.81%	22.67%

Exploring Pairwise Politeness vs. Stance Similarity Given Social Communities

Table 4 shows stance similarity and interaction politeness characteristics at the level of communities. When one considers edges where both the users are members of the same community, a higher fraction of pairs are similar, but the average politeness is lower, as compared to where the two users belong to two different communities. Note that, all of these user pairs are familiar, and this observation holds true at a community (vs. non-community) level rather than at the level of specific familiarities. Again, a manual inspection into the dataset suggests less formalness within communities (which largely comprises of familiar individuals).

Summary of Observations

In summary, we observe that, similar people, when they are also familiar tend to be more polite, compared to similar people who are not familiar. However, in the case of dissimilar people, when they are not familiar they tend to be more formal (hence, polite). In communities too, non-familiar users, who are dissimilar, are more polite.

Table 3. An intuition of the raw results. Examples have been randomly picked from the dataset and results, covering cases where similarity and politeness values are high and low, as well as the user pairs are familiar vs. not. S ← similar, NS ← non-similar, P ← polite, NP ← not polite.

Sample of hashtags used by User 1	Sample of hashtags used by User 2	Are they familiar?	Similarity score	Politeness score	Example Type
followfriday, computerworld, Twitter, feedly	navigon, mustsee, vgfail, dilbert, humor	Yes	−0.48538217	0.302404184	NS, NP
unfollowfriday, byebye	byebye, iranelection	Yes	0.32039406	0.29075952	S, NP
140conf, journchat	140conf	Yes	0.875587667	0.888408196	S, P
ff, fcxp09	iranelection, followfriday	Yes	−0.375685965	0.929627317	NS, P
physics, transhumanism, science, quantam	fun, quote, art	No	−0.265319422	0.373339869	NS, NP
grn, followfriday	webdesign, abuzz, grn, followfriday	No	0.816427949	0.259325957	S, NP
squarespace, foursquare, confedcup	squarespace	No	0.723357763	0.773153153	S, P
iranelection	oldmusicwednesday	No	−0.536816617	0.619809616	NS, P

Table 4. Stance similarity and interaction politeness characteristics at a community level. *Percentage similar* denotes the percentage of similar user pairs, and *Percentage polite* denotes the percentage of interactions that are polite in nature. Politeness threshold 0.67 ("highly polite") and similarity threshold 0.33 are used.

Type of the edges	Number of edges	Are user pairs similar?	Are user pairs polite?
Edges within communities	3,219,879	35.09%	6.84%
Edges across communities	1,757,618	34.68%	8.12%

5 Conclusion

In this work, we empirically assessed the impact of familiarity and similarity in social information and interaction dynamics. We used stance - the sentiment of users with respect to a certain given topic - as the anchor of similarity, wherein, hashtags were used as individual topics. This is a novel problem area, so far unexplored in the literature, given that the stance detection problem on Twitter has emerged as recently as 2016. Specifically, we attempted to answer two key questions. First, whether higher similarity of stance across different topics amongst users, leads to smoother (more polite) conversation, and lower similarity of stances tend to lead to higher frictions (less politeness). Second, whether individuals belonging to the same implicit community (formed by social friendships that indicate explicit familiarity), tend to have similar stances towards hashtags. The second question essentially explores a fundamental question about social networks: whether or not homophily holds at a finer grain (stances) rather than the coarser grain (sentiments, direct interests or explicit social group memberships). Our work is applicable in social information analysis and social network based marketing optimization applications, in practice.

References

1. Aral, S., Muchnik, L., Sundararajan, A.: Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. PNAS **106**(51), 21544–21549 (2009)
2. Bengio, Y., Ducharme, R., Vincent, P., Jauvin, C.: A neural probabilistic language model. J. Mach. Learn. Res. **3**, 1137–1155 (2003)
3. Bhattacharya, P., Goyal, P., Sarkar, S.: Query translation for cross-language information retrieval using multilingual word clusters. WSSANLP **2016**, 142 (2016)
4. Blei, D.M., Ng, A.Y., Jordan, M.I.: Latent dirichlet allocation. J. Mach. Learn. Res. **3**, 993–1022 (2003)
5. Blondel, V.D., Guillaume, J.L., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. J. Stat. Mech. Theory Experiment **2008**(10), P10008 (2008)
6. Boltuzic, F., Karan, M., Alagic, D., Šnajder, J.: Takelab at semeval-2016 task 6: stance classification in tweets using a genetic algorithm based ensemble. In: Proceedings of SemEval, pp. 464–468 (2016)
7. Cataldi, M., Di Caro, L., Schifanella, C.: Emerging topic detection on twitter based on temporal and social terms evaluation. In: Tenth International Workshop on Multimedia Data Mining, p. 4. ACM (2010)
8. Cataldi, M., Schifanella, C., Candan, K.S., Sapino, M.L., Di Caro, L.: Cosena: a context-based search and navigation system. In: International Conference on Management of Emergent Digital EcoSystems, p. 33. ACM (2009)
9. Cunha, E., Magno, G., Comarela, G., Almeida, V., Gonçalves, M.A., Benevenuto, F.: Analyzing the dynamic evolution of hashtags on twitter: a language-based approach. In: Proceedings of the Workshop on Languages in Social Media, pp. 58–65. ACL (2011)
10. Danescu-Niculescu-Mizil, C., Sudhof, M., Jurafsky, D., Leskovec, J., Potts, C.: A computational approach to politeness with application to social factors. arXiv preprint arXiv:1306.6078 (2013)
11. De Choudhury, M., Sundaram, H., John, A., Seligmann, D.D., Kelliher, A.: "birds of a feather": does user homophily impact information diffusion in social media? arXiv preprint arXiv:1006.1702 (2010)
12. Dey, K., Shrivastava, R., Kaushik, S.: A paraphrase and semantic similarity detection system for user generated short-text content on microblogs. In: COLING, pp. 2880–2890 (2016)
13. Dey, K., Shrivastava, R., Kaushik, S.: Twitter stance detection - a subjectivity and sentiment polarity inspired two-phase approach. In: SENTIRE Workshop, ICDM (2017)
14. Dey, K., Shrivastava, R., Kaushik, S., Subramaniam, L.V.: Emtagger: a word embedding based novel method for hashtag recommendation on twitter. In: ACUMEN Workshop, ICDM (2017)
15. Halberstam, Y., Knight, B.: Homophily, group size, and the diffusion of political information in social networks: evidence from twitter. J. Public Econ. **143**, 73–88 (2016)

16. Kwak, H., Lee, C., Park, H., Moon, S.: What is twitter, a social network or a news media? In: WWW, pp. 591–600. ACM (2010)
17. Lau, J.H., Collier, N., Baldwin, T.: On-line trend analysis with topic models:\# twitter trends detection topic model online. In: COLING, pp. 1519–1534 (2012)
18. Mathioudakis, M., Koudas, N.: Twittermonitor: trend detection over the twitter stream. In: SIGMOD, pp. 1155–1158. ACM (2010)
19. McPherson, M., Smith-Lovin, L., Cook, J.M.: Birds of a feather: homophily in social networks. Annu. Rev. Sociol. **27**(1), 415–444 (2001)
20. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. arXiv preprint arXiv:1301.3781 (2013)
21. Mishra, A., Dey, K., Bhattacharyya, P.: Learning cognitive features from gaze data for sentiment and sarcasm classification using convolutional neural network. In: ACL (vol. 1: Long Papers), vol. 1, pp. 377–387 (2017)
22. Mohammad, S.M., Kiritchenko, S., Sobhani, P., Zhu, X., Cherry, C.: Semeval-2016 task 6: detecting stance in tweets. In: SemEval, vol. 16 (2016)
23. Mohammad, S.M., Sobhani, P., Kiritchenko, S.: Stance and sentiment in tweets. arXiv preprint arXiv:1605.01655 (2016)
24. Musto, C., Semeraro, G., De Gemmis, M., Lops, P.: Word embedding techniques for content-based recommender systems: an empirical evaluation. In: RecSys Posters (2015)
25. Newman, M.E.: Modularity and community structure in networks. Proc. Natl. Acad. Sci. **103**(23), 8577–8582 (2006)
26. Nobata, C., Tetreault, J., Thomas, A., Mehdad, Y., Chang, Y.: Abusive language detection in online user content. In: WWW, pp. 145–153 (2016)
27. Papegnies, E., Labatut, V., Dufour, R., Linares, G.: Impact of content features for automatic online abuse detection. arXiv preprint arXiv:1704.03289 (2017)
28. Pennington, J., Socher, R., Manning, C.D.: Glove: global vectors for word representation. EMNLP **14**, 1532–1543 (2014)
29. Rodríguez, M.A., Egenhofer, M.J.: Determining semantic similarity among entity classes from different ontologies. IEEE Trans. Knowl. Data Eng. **15**(2), 442–456 (2003)
30. Šćepanović, S., Mishkovski, I., Gonçalves, B., Nguyen, T.H., Hui, P.: Semantic homophily in online communication: evidence from twitter. Online Soc. Netw. Media **2**, 1–18 (2017)
31. Schmidt, A., Wiegand, M.: A survey on hate speech detection using natural language processing. In: Workshop on Natural Language Processing for Social Media. ACL, Valencia, pp. 1–10 (2017)
32. Sidorov, G., Gelbukh, A., Gómez-Adorno, H., Pinto, D.: Soft similarity and soft cosine measure: Similarity of features in vector space model. Computación y Sistemas **18**(3), 491–504 (2014)
33. Stieglitz, S., Dang-Xuan, L.: Emotions and information diffusion in social mediasentiment of microblogs and sharing behavior. J. Manag. Inf. Syst. **29**(4), 217–248 (2013)
34. Vijayaraghavan, P., Sysoev, I., Vosoughi, S., Roy, D.: Deepstance at semeval-2016 task 6: detecting stance in tweets using character and word-level cnns. arXiv preprint arXiv:1606.05694 (2016)
35. Weston, J., Chopra, S., Adams, K.: # tagspace: semantic embeddings from hashtags (2014)
36. Yang, J., Leskovec, J.: Patterns of temporal variation in online media. In: WSDM, pp. 177–186. ACM (2011)

37. Zarrella, G., Marsh, A.: Mitre at semeval-2016 task 6: transfer learning for stance detection. arXiv preprint arXiv:1606.03784 (2016)
38. Zhang, Z., Lan, M.: Ecnu at semeval-2016 task 6: relevant or not? supportive or not? a two-step learning system for automatic detecting stance in tweets. SemEval, pp. 451–457 (2016)
39. Zhao, J., Lan, M., Tian, J.F.: Ecnu: using traditional similarity measurements and word embedding for semantic textual similarity estimation. In: SemEval, pp. 117–122 (2015)