

Improving Twitter Community Detection through Contextual Sentiment Analysis of Tweets

Alron Jan Lam

College of Computer Studies
De La Salle University
2401 Taft Avenue, Malate, Manila, Philippines
alron_lam@dlsu.edu.ph

Abstract

Works on Twitter community detection have yielded new ways to extract valuable insights from social media. Through this technique, Twitter users can be grouped into different types of communities such as those who have the same interests, those who interact a lot, or those who have similar sentiments about certain topics. Computationally, information is represented as a graph, and community detection is the problem of partitioning the graph such that each community is more densely connected to each other than to the rest of the network. It has been shown that incorporating sentiment analysis can improve community detection when looking for sentiment-based communities. However, such works only perform sentiment analysis in isolation without considering the tweet's various contextual information. Examples of these contextual information are social network structure, and conversational, author, and topic contexts. Disregarding these information poses a problem because at times, context is needed to clearly infer the sentiment of a tweet. Thus, this research aims to improve detection of sentiment-based communities on Twitter by performing contextual sentiment analysis.

1 Introduction

Twitter, as a micro-blogging platform, has become an avenue for people to voice out their opinions online. This gives concerned entities, like policy-makers or brand managers, the chance to hear people out in an unprecedented way. However, to effectively utilize this source of information, the

massive amount of tweets must first be processed to be more easily understood (Kavanaugh et al., 2012).

One such way to achieve this is through community detection, which is a domain-independent, graph-based problem that can be applied to many different disciplines including social media analysis. Its definition is that it is the problem of looking for groups of vertices that are more densely connected to each other than to the rest of the graph (Papadopoulos et al., 2012; Tang and Liu, 2010). Hence, due to its domain-independence, setting up the input graph properly according to the desired application is important (Darmon et al., 2014). When applied to Twitter, a wide array of communities can be found, such as communities of users with similar interests (Lim and Datta, 2012b; Lim and Datta, 2012a), communities of users who interact frequently (Amor et al., 2015), communities of geographically-nearby users (Bakillah et al., 2015), or communities of users with similar sentiments towards a certain topic (Cao et al., 2015), among many other possibilities.

Finding these communities can yield insights like (1) the different kinds of conversations going on, (2) who participates in them, and (3) what they are talking about. These kinds of insights could be valuable to entities such as policy-makers (Amor et al., 2015).

However, most works on Twitter community detection are focused on finding communities within social networks based on explicit relationships between users, such as following-follower relationships (Java et al., 2007), or mention/re-tweet relationships (Pearce et al., 2014; Zhang et al., 2012). An often-overlooked source of information are the actual contents of the users' tweets. In some cases, this may not be important. But when looking for communities of users who share similar sentiments, this could potentially improve

community detection (Deitrick and Hu, 2013).

The work of Deitrick and Hu (2013) utilized sentiment analysis to improve community detection. In addition to the usual graph edges representing follower-following relations, they added more edge weights between users who expressed the same sentiments towards the same topic. That is, whenever two users tweet with the same sentiment polarity (positive or negative) containing the same hashtag (treated as the topic), their edge weight is incremented, indicating a stronger relationship between these users. They showed that doing this technique improves community detection according to the modularity score, and this in turn, facilitates better sentiment analysis.

However, works like that of Deitrick and Hu (2013) perform sentiment analysis of a tweet in an isolated manner. That is, various contextual information available for a tweet are totally disregarded. Examples of these include conversational (tweets preceding the target tweet in the conversation thread), author (tweets recently posted by the author preceding the target tweet) and topic (recent tweets about the same topic posted before the target tweet) context (Ren et al., 2016; Vanzo et al., 2014). Another example of contextual information is social network structure, wherein connections between users help determine sentiment polarities by utilizing social theories such as balance theory ("an enemy of my enemy is my friend") and homophily ("birds of the same feather flock together") (West et al., 2014; Tan et al., 2011). The aforementioned studies have shown that incorporating contextual information can improve sentiment analysis.

Thus, in looking for sentiment-based Twitter communities wherein there are stronger connections between users having similar sentiments, it may be beneficial to take a contextual approach to sentiment analysis.

2 Twitter Community Detection

This section compares and contrasts different works on Twitter community detection. Table 1 is a summary of all reviewed works and a comparison of these works in terms of desired community type and edge construction scheme. The community detection algorithms and evaluation metrics used by these works are also discussed in this section.

It can be observed in Table 1 that different re-

Year	Authors	Community Types	Edge Construction
2013	Deitrick & Wu	Social Network-based, Sentiment-based	Based on follow, mention, and re-tweet relationships, and on tweets having the same hashtags and same sentiment polarity
2014	Darmon et al.	Interaction-based, Topic-based	Based on mention and re-tweet relationships, and on tweets having the same hashtags
2015	Bakillah et al.	Interaction-based, Topic-based	Based on follow, mention, and re-tweet relationships, and on tweets having the same URLs and similar tweet content
2015	Amor et al.	Social Network-based, Interaction-based	Based on follow, mention, and re-tweet relationships
2015	Cao et al.	Sentiment-based	Based on difference between users' sentiment trends over time

Table 1: Twitter Community Detection Works and their Desired Community Types, and Edge Construction Scheme

search works aimed to identify different types of communities. These communities are: (1) social network-based (who are in the same social groups?), (2) interaction-based (who frequently interact with each other through re-tweets and mentions?), (3) topic-based (who talk about the same things?), and (4) sentiment-based (who feel the same way about certain topics?).

It is important to note the different types of communities the works aimed to extract, because they largely influence how the community detection problem is set-up. Since community detection is a domain-independent problem that applies to any graph structure, performing Twitter community detection requires that relevant information be represented appropriately through the vertices and edges in a graph. With the input graph laid out, researchers can then select appropriate community detection algorithms to be used. Lastly, researchers can then choose appropriate metrics for evaluating their approach. These three aspects are discussed in more detail in the following subsections.

2.1 Edge Construction

As previously mentioned, the desired community types largely influence the representation of information as a graph, consisting of vertices and edges. Although in theory this might not always be the case, most reviewed works used vertices to represent users, and consequently, edges to represent relationships between the users. The works only differed in terms of what kind of relationships the edges represented.

For works that aimed to identify social groups, that is, communities of people who generally followed each other within their own communities, edges have generally been used to represent a “follow” relationship (Amor et al., 2015; Lim and Datta, 2012b; Lim and Datta, 2012a). On Twitter, users “follow” other users to “subscribe” to their tweets and be able to see them as they get published. It is important to note that the “follow” relationship is not necessarily two-way. That is, if user A follows user B, it is possible that user B does not follow user A. Given the explicit “follow” relationships on Twitter, some works have represented “follow” networks on Twitter in a straightforward manner by using directed edges to represent these “follow” relationships. As discussed earlier, the goal of these works is usu-

ally to find social groups or cliques within the graph that represent social circles (since people who are friends in real life tend to follow each other) or people with similar interests (since people also follow users they may not know personally, but whom they are interested in). This type of edge construction, since it is straightforward and based on explicit “follow” relationships, is usually used in combination with other edge construction schemes, wherein the “follow” relationships determine the existence of edges, while other information are used to increment edge weights (Deitrick and Hu, 2013).

For works that aimed to identify communities of people that interacted with each other frequently, main relationships involved are the “re-tweet” and “mention” relationships (Amor et al., 2015; Darmon et al., 2014; Bakillah et al., 2015). When user A re-tweets a tweet by user B, user A is essentially re-publishing the said tweet and propagating it to his/her own followers. On the other hand, a mention happens when user A tags user B in a tweet. This either happens when user A simply wants to call user B’s attention to his/her tweet, or when user A is replying to one of user B’s tweets. Note that it is possible for users to mention multiple users in a single tweet. Having said that, mentions or re-tweets between users have been used to either increment existing edge weights (Amor et al., 2015; Darmon et al., 2014), or to establish the existence of new edges altogether (Bakillah et al., 2015) in works that sought interaction-based communities. An example of such a work is that of Amor et al. (2015), where they found what they called conversation communities: groups of conversations with people talking about a topic of interest.

For works that aimed to identify communities of people who were interested in similar topics (Bakillah et al., 2015; Darmon et al., 2014; Deitrick and Hu, 2013), hashtags have been used to establish or strengthen edges between users. Hashtags are a way for Twitter users to tag their tweet as talking about a certain topic. These are arbitrarily defined and purely user-generated tags in the form ‘#hashtag’. Users sometimes tend to copy other hashtags instead of creating their own, resulting into popular, trending hashtags. With this in mind, the idea of topic-based community detection is to look for communities of users who talk about similar topics identified through the hash-

tags. For example, if user A tweets with the hashtag '#ClimateChange', and user B also tweets with the same hashtag, then either more weight is added to an existing edge between them (Darmon et al., 2014; Deitrick and Hu, 2013), or a new edge between the users is created on this basis (Bakillah et al., 2015).

For works that aimed to identify communities of people sharing similar sentiments, the idea is to establish stronger relationships between users who feel the same sentiment polarity toward the same topic (identified through the hashtag) (Cao et al., 2015; Deitrick and Hu, 2013). For example, if user A tweets "We are in a bad situation. #ClimateChange" and user B tweets "Our world is dying. #ClimateChange", then user A and user B's edge should be added more weight because their tweets both express negative sentiments about climate change (Deitrick and Hu, 2013).

In summary, different community types warrant different edge construction schemes. However, it is important to note that works on Twitter community detection do not necessarily utilize just one of the aforementioned schemes. Rather, researchers oftentimes experiment with and combine different edge construction and weighting schemes to see which configuration produces the best output (Bakillah et al., 2015; Deitrick and Hu, 2013).

2.2 Algorithms

The reviewed works on community detection have used a variety of algorithms, each being appropriate to different scenarios or constraints. For example, some of the algorithms can handle directed and weighted graphs (Xie, 2012; Rosvall and Bergstrom, 2008; Lancichinetti et al., 2011), while some can detect overlapping communities (Xie, 2012; Lancichinetti et al., 2011), while some execute relatively quickly for large graphs (Xie, 2012; Rosvall and Bergstrom, 2008). These are examples of factors that the researchers took into consideration when choosing their algorithms. A more detailed discussion of each work follows.

Deitrick and Hu (2013) chose the Speaker Listener Propagation Algorithm, or SLPA, (Xie, 2012) and the Infomap algorithm (Rosvall and Bergstrom, 2008) for community detection as they both work with weighted and directed graphs, and they both execute relatively quickly on large graphs. In addition, SLPA can identify overlapping communities.

Darmon et al. (2014) chose the Order Statistics Local Optimization Method, or OSLOM, (Lancichinetti et al., 2011) for community detection because of its ability to work with weighted and directed graphs, and its ability to identify overlapping communities.

Bakillah et al. (2015) chose the Fast-Greedy Optimization of Modularity, or FGM, (Clauset et al., 2004) for its ability to handle complex social graphs from Twitter, and the Varied Density-Based Spatial Clustering of Applications with Noise, or VDBSCAN, (Liu et al., 2007) for its ability to obtain spatial clusters at certain points in time.

Amor et al. (2015) chose the Markov Stability (Delvenne et al., 2010) due to its mechanism of modeling information flow. The primary goal of their research was to understand the Twitter discussion on the care.data program in terms of information flow and the roles that Twitter users play. Hence, their selection of Markov Stability fits their goals.

Lastly, Cao et al. (2015) chose to use Hierarchical Agglomerative Clustering (Jain and Dubes, 1988) based on sentiment distance. Since they were focused on looking for communities with similar sentiments, the clustering method is appropriate for this task.

2.3 Evaluation

To evaluate their approaches, researchers of related works have used quantitative and/or qualitative analysis. Quantitative analysis usually entails optimizing some metric, like the well-known modularity score (Newman, 2006), which indicates how well-formed the communities are as opposed to randomly generated communities. Other works have also performed experiments in which they pre-determined the communities beforehand, treating the community detection problem as a "classification" problem of placing vertices in their proper communities. As such, these works have used precision and recall to evaluate their approach (Bakillah et al., 2015). However, the manual pre-determination of communities beforehand can be a difficult task, so this kind of evaluation methodology is not too popular, making metric optimization as the more common evaluation approach. Exact numerical results of these studies are not discussed here because direct comparison of results is not appropriate due to differences in datasets and types of communities being detected.

Year	Authors	Level of Sentiment Analysis	Contextual Information Used
2016	Ren et al.	Document-level (tweet)	Conversational, Author, Topic Context
2014	Vanzo et al.	Document-level (tweet)	Conversational Context
2014	West et al.	User-level (towards another user)	Social Network Structure
2011	Tan et al.	User-level (towards a topic)	Social Network Structure

Table 2: Contextual Sentiment Analysis Works and their Levels of Sentiment Analysis and Contextual Information Used

On the other hand, researchers also use qualitative analyses in the form of case studies. Usually, this comes in the form of a discussion on the insights acquired from the community detection approach. For example, Amor et al. (2015) discussed in their work how their approach was able to reveal insights into who were concerned about the care.data program in the UK (political activists, media, UK healthcare professionals, and US healthcare professionals) and what they were concerned about (data privacy, impact on patient welfare, etc). Other works like that of Cao et al. (2015) also involved interviewing domain experts and asking them to evaluate whether community detection results would be useful to them or others in the field whose tasks involve analyzing social media data.

3 Contextual Sentiment Analysis

This section compares and contrasts different works on Contextual Sentiment Analysis. Shown in Table 2 is a summary of reviewed works and a comparison of these works in terms of level of sentiment analysis and context types considered. The algorithms and evaluation metrics used by these works are also discussed in the section.

3.1 Sentiment Analysis Types

It can be seen in Table 2 that for the reviewed works, there are two levels of sentiment analy-

sis: document-level and person-level. A document is essentially a collection of sentences. In the case of Ren et al. (2016), Vanzo et al. (2014), and Tan et al. (2011), having Twitter as the domain, a document refers to a tweet. While in the case of West et al. (2014) with Wikipedia discussions and US Congress speeches as the domain, a document refers to a person’s post and speech respectively. Document-level sentiment analysis usually involves utilizing lexical information found in the text.

On the other hand, person-level sentiment analysis focuses on determining the overall sentiment of a person towards a particular person on topic, as opposed to focusing on each individual document a person generates. For example, say it is desirable to determine user A’s sentiment towards Obama based on his/her tweets. Person-level sentiment analysis would then require consideration of all user A’s tweets about Obama, instead of just determining the conveyed sentiment in each tweet. For most related works, document-level sentiment analysis is performed as a sub-task of person-level sentiment analysis (West et al., 2014; Tan et al., 2011).

3.2 Context Types

In addition to using textual information for sentiment analysis, the reviewed works utilized a variety of contextual information. The principle is that these provide more knowledge needed to perform more accurate sentiment analysis.

For document-level sentiment analysis (tweet-level in this case), context types used by Ren et al. (2016) and Vanzo et al. (2014) are conversational, author and topic. Having Twitter as the domain, conversational context was defined as the most recent l tweets preceeding a target tweet in the conversation it belongs to. Author context was defined by Ren et al. (2016) as the most recent l tweets posted by a user before the target tweet. Lastly, topic context was defined by Ren et al. (2016) as the most recent l tweets posted before the target tweet that shares at least one hashtag with the target tweet. The rationale is that the textual information found in a single tweet may be ambiguous, and thus, insufficient to clearly determine its sentiment polarity. Therefore, taking into account the aforementioned contexts can fill in the said gap.

For person-level sentiment analysis, social net-

work structure has been used by West et al. (2014) and Tan et al. (2011) as contextual information. These works rely on theories about social behavior such as balance theory ("an enemy of my enemy is my friend") and homophily ("birds of the same feather flock together") to complement document-level sentiment analysis based on the document text. The idea is that information gained from people's connections or interactions can help determine a person's sentiment towards a topic (Tan et al., 2011) or another user (West et al., 2014).

3.3 Methodology

Ren et al. (2016), with document-level sentiment analysis as the goal, represented words found in the target tweet and in contextual tweets (conversational, author, and topic, as explained in the previous sub-section) through word embeddings. They then train a convolutional neural network to classify the target tweet's sentiment polarity given these input features.

Vanzo et al. (2014), with document-level sentiment analysis as the goal, use a Markovian formulation of the Support Vector Machine model to classify a target tweet given the preceding tweets in its conversational context. They represent tweets through bag of words, a distributed lexical semantic model, a user sentiment inclination profile, and various combinations of these three.

West et al. (2014), with person-level sentiment analysis (towards another person) as the goal, use a scheme they call "triangle balance" in which they minimize a cost function that applies penalties for going against the sentiment model and for going against the social theories they used. The setting of cost parameters was done through machine learning.

Tan et al. (2011), with person-level sentiment analysis (towards a topic) as the goal, use a factor-graph model for estimating the probability of each polarity for a given person. They experiment on learning and no-learning approaches in setting the necessary parameters.

Since the determination of sentiment polarities is generally a classification problem, most of the reviewed works evaluated their results through metrics common to classification tasks like precision, recall, F-measure, accuracy, and ROC.

4 Conclusion

Based on the review of related works, it can be seen that the desired community types largely dictate the edge construction scheme used in the input graphs. Furthermore, it has been shown that using sentiment analysis to modify edge weights when performing community detection can improve the detection of sentiment-based communities (Deitrick and Hu, 2013). The idea is that users who feel the same about a particular topic should have a stronger connection.

However, one possible improvement over the work of (Deitrick and Hu, 2013) is to perform contextual sentiment analysis. This is because various contextual information, such as conversational, author, and topic context, along with social network structure, have been shown to improve sentiment analysis (Ren et al., 2016; Vanzo et al., 2014; West et al., 2014; Tan et al., 2011). The assumption is that the improvement in sentiment analysis will improve the modification of edge weights (and therefore, the representation of connection between users) and consequently, improve sentiment-based community detection. Evaluation can be through quantitative analysis by using well-known metrics in community detection, such as modularity, or through qualitative analysis by performing case studies. Analysis of the results can provide insight on which contextual information provide the most improvement in the task of sentiment-based community detection on Twitter.

References

- B Amor, S Vuik, Ryan Callahan, Ara Darzi, Sophia N Yaliraki, and Mauricio Barahona. 2015. Community detection and role identification in directed networks: understanding the twitter network of the care. data debate. *arXiv preprint arXiv:1508.03165*.
- Mohamed Bakillah, Ren-Yu Li, and Steve HL Liang. 2015. Geo-located community detection in twitter with enhanced fast-greedy optimization of modularity: the case study of typhoon haiyan. *International Journal of Geographical Information Science*, 29(2):258–279.
- Nan Cao, Lu Lu, Yu-Ru Lin, Fei Wang, and Zhen Wen. 2015. Socialhelix: visual analysis of sentiment divergence in social media. *Journal of Visualization*, 18(2):221–235.
- Aaron Clauset, Mark EJ Newman, and Cristopher Moore. 2004. Finding community structure in very large networks. *Physical review E*, 70(6):066111.

- David Darmon, Elisa Omodei, and Joshua Garland. 2014. Followers are not enough: A question-oriented approach to community detection in online social networks. *arXiv preprint arXiv:1404.0300*.
- William Deitrick and Wei Hu. 2013. Mutually enhancing community detection and sentiment analysis on twitter networks.
- J-C Delvenne, Sophia N Yaliraki, and Mauricio Barahona. 2010. Stability of graph communities across time scales. *Proceedings of the National Academy of Sciences*, 107(29):12755–12760.
- Anil K Jain and Richard C Dubes. 1988. *Algorithms for clustering data*. Prentice-Hall, Inc.
- Akshay Java, Xiaodan Song, Tim Finin, and Belle Tseng. 2007. Why we twitter: understanding microblogging usage and communities. In *Proceedings of the 9th WebKDD and 1st SNA-KDD 2007 workshop on Web mining and social network analysis*, pages 56–65. ACM.
- Andrea L Kavanaugh, Edward A Fox, Steven D Sheetz, Seungwon Yang, Lin Tzy Li, Donald J Shoemaker, Apostol Natsev, and Lexing Xie. 2012. Social media use by government: From the routine to the critical. *Government Information Quarterly*, 29(4):480–491.
- Andrea Lancichinetti, Filippo Radicchi, José J Ramasco, and Santo Fortunato. 2011. Finding statistically significant communities in networks. *PloS one*, 6(4):e18961.
- Kwan Hui Lim and Amitava Datta. 2012a. Finding twitter communities with common interests using following links of celebrities. In *Proceedings of the 3rd international workshop on Modeling social media*, pages 25–32. ACM.
- Kwan Hui Lim and Amitava Datta. 2012b. Following the follower: detecting communities with common interests on twitter. In *Proceedings of the 23rd ACM conference on Hypertext and social media*, pages 317–318. ACM.
- Peng Liu, Dong Zhou, and Naijun Wu. 2007. Vdb-scan: varied density based spatial clustering of applications with noise. In *Service Systems and Service Management, 2007 International Conference on*, pages 1–4. IEEE.
- Mark EJ Newman. 2006. Modularity and community structure in networks. *Proceedings of the national academy of sciences*, 103(23):8577–8582.
- Symeon Papadopoulos, Yiannis Kompatsiaris, Athena Vakali, and Ploutarchos Spyridonos. 2012. Community detection in social media. *Data Mining and Knowledge Discovery*, 24(3):515–554.
- Warren Pearce, Kim Holmberg, Iina Hellsten, and Brigitte Nerlich. 2014. Climate change on twitter: Topics, communities and conversations about the 2013 ipcc working group 1 report. *PloS one*, 9(4):e94785.
- Yafeng Ren, Yue Zhang, Meishan Zhang, and Donghong Ji. 2016. Context-sensitive twitter sentiment classification using neural network. AAAI.
- Martin Rosvall and Carl T Bergstrom. 2008. Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences*, 105(4):1118–1123.
- Chenhao Tan, Lillian Lee, Jie Tang, Long Jiang, Ming Zhou, and Ping Li. 2011. User-level sentiment analysis incorporating social networks. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1397–1405. ACM.
- Lei Tang and Huan Liu. 2010. Community detection and mining in social media. *Synthesis Lectures on Data Mining and Knowledge Discovery*, 2(1):1–137.
- Andrea Vanzo, Danilo Croce, and Roberto Basili. 2014. A context-based model for sentiment analysis in twitter. In *COLING*, pages 2345–2354.
- Robert West, Hristo S Paskov, Jure Leskovec, and Christopher Potts. 2014. Exploiting social network structure for person-to-person sentiment analysis. *arXiv preprint arXiv:1409.2450*.
- Jierui Xie. 2012. *Agent-based dynamics models for opinion spreading and community detection in large-scale social networks*. Ph.D. thesis, Rensselaer Polytechnic Institute.
- Yang Zhang, Yao Wu, and Qing Yang. 2012. Community discovery in twitter based on user interests. *Journal of Computational Information Systems*, 8(3):991–1000.