Are You a Compatible User?
Compatibility of a Microblog User with a News Article

Sang-Pil Kim and Byung Suk Lee

Department of Computer Science, University of Vermont, Burlington VT 05405, USA
skim21@uvm.edu, bslee@uvm.edu

Abstract. A novel concept of compatibility between a news article and an online microblog user is introduced, and a framework embodying the concept is proposed. The framework currently proposes to match two factors – user's interest and user's sentiment as reflected in the user's microblog texts – to determine the compatibility. Using Twitter as an example, the framework is instantiated using the RAKE algorithm for topic keyword (for interest) matching and the VADER model-based sentiment scoring algorithm for sentiment matching. Gold standard tests show that considering both interest match and sentiment match improves the accuracy of compatibility decision significantly and that filtering topic keywords based on co-occurrence semantics helps to disambiguate the user's sentiment match, hence the compatibility decision.

Keywords: online news media; online social media; user compatibility; user interest match; user sentiment match

1 Introduction

News media went online from the inception of online social media, and now they are one of the major driving forces behind the interplay among content providers, consumers, and commentators, involving millions of users in complex real-time social dynamics. Online social media is a domain where a large number of users interact, share, and disseminate information freely, and there has been an increasing interest in identifying users who appear to be expressing opinions compatible with online media news. Users found as such can be targets of marketing in sales, candidates of polls in politics, or potential donors in fund raising, to mention a few.

In this regard, a larger goal of our work is to efficiently and accurately find users that are most compatible with a certain news article. The specific work presented in this paper focuses on one issue critical to achieving the objective, that is, profiling whether a given user (who is an online news reader) is compatible with a news article posted in online social media. We have chosen Twitter as the online social media because of its ability to disseminate information rapidly across extensive user base and also because of the challenges stemming from its being *micro*blogging (i.e., no more than 140 characters).

The notion of compatibility between news and users has no established definition yet, and so we mean to start the discussion and propose an initial set of factors for judging the compatibility between a news article and a Twitter user. In this paper, we focus on user's interest and user's sentiment as the two key baseline factors. These two require analyzing microblogs to see how well they match with the interest and sentiment reflected in a news article, and are founded upon bodies of work in the respective areas of topic mining (e.g. [3][7][24]) and sentiment analysis (e.g., [13][16]). Thus, the protocol we currently use for the compatibility is a two-step approach – perform the *user interest matching* first and then the *user sentiment matching* next.

One lesson learned during the work was due to the sparsity problem, which is inherent in tweets because of their limit on the length and is worsened by the recently emerging flimsy tweeting behavior of users. That is, tweet users do not write much in their messages – many times they simply retweet other tweets or include links to other media materials (e.g., longer text, images, audio, video), without adding substantial contents of their own. This phenomenon – well known in the community (e.g., [15]) – drove us to be inclusive in identifying a user's tweets that are relevant to (hence match) the topic of a news article and judging if the relevant tweets of the user show the same sentiment as the news article. In addition, we leveraged the bootstrapping technique (i.e., bundling up multiple tweet texts together into a longer text) as used in other work (e.g., [5][18][23]) to overcome the same problems of tweets.

Another lesson learned is that, while using multiple topic keywords helps with determining the user's interest match, it brings ambiguities in the subsequent sentiment analysis. The reason is that the selected keywords have no bearing on the sentiment of the user's tweets, and, therefore, often lead to opposite sentiments, thus cancelling each other in determining the polarity (i.e., positive or negative) of the user's sentiment. Our finding in this regard is that taking advantage of the keywords' co-occurrence semantics to filter out some of the keywords helps to resolve this problem greatly.

Evaluations were done with regard to the two factors (i.e., user interest and user sentiment). A real tweet dataset, collected from Tweet User API was used for experiments. Ground truth was constructed with news articles and Tweet users selected manually based on their contents' relevance and sentiments. Algorithms that consider either of the two factors separately were used as the baselines. The gold standard test using the ground truth showed that the average f-score achieved by our algorithm (with keyword filtering) was higher than those achieved by the interest-match-only algorithm and the sentiment-match-only algorithm by 2.07 and 2.93 times, respectively. The test also showed that filtering out keywords that have weak co-occurrence relationships improved the f-score 1.43 times.

We claim the merit of this paper in being the first to introduce the notion of compatibility between a news article and online social media users, which finds a lot of real world applications. Additionally, this paper introduces a framework and presents an implementation by leveraging proven state of the art algorithms, which not only demonstrates the feasibility of assessing the compatibility but also provides a baseline for further research on this topic.

The remainder of the paper is organized as follows. Section 2 discusses related work. Section 3 describes the proposed method for judging if a user is compatible with a news article. Section 4 presents the experiments and results. Section 5 concludes the paper with a summary and an outline of further work.

2 Related work

To the best of our knowledge, there is no prior work done by others to consider both the interest and the sentiment of a user in light of those reflected in a new article.

Some research focused on user's interest as reflected from their tweets – to determine the user's interests themselves [10][14] or to classify users based on their interests [1][2][11]. Kapanipathi et al. [10] determined the topic of user's interests by generating a Wikipedia "hierarchy" and determining the topic from the user's tweets from the hierarchy. Michelson and Macskassy [14] built a "topic profile" from multiple users and identified the topic of interest for a given user from it. They also used Wikipedia to disambiguate the concepts mentioned in tweets. Alvarez-Melis and Saveski [1] proposed a new "pooling technique" by which tweets exchanged between users are grouped based on topic modeling. Lim and Datta [12] picked a user that had more than 10,000 followers and classified the followers to categorize their interests into 15 topics using the Wikipedia hierarchy. Campbell et al. [2] proposed a method to classify users based on both tweet content and contextual information (e.g., retweets, mentions, co-occurrences).

Some research focused on sentiment to classify users based on the sentiment reflected in their tweets [6][21]. Gutierrez and Poblete [6] clustered users based on their sentiment polarity trace to generate individual profiles for different concepts, and studied the characteristics of the clusters. Tan et el. [21] proposed to take advantage of social relationships (e.g., followship, homophily, approval) between users to improve the sentiment-based user classification.

Some research used both interest and sentiment in their work [4][17]. Chen and Mirisaee's work [4] on topic-based sentiment analysis has something common with ours in that they built topic-dependent models – one driven by a target keyword and one driven by a group of topic-related tweet terms – and used them in sentiment analysis. Their work, however, is not concerning users at all. Pennacchiotti and Popescu [17] developed a generic model for user classification with a larger scope based on a comprehensive set of key features such as user's attributes (e.g., name, location), user's tweet behavior (e.g., frequencies of different types of tweets), user's tweet contents (e.g., keywords, hashtags, topic word, sentiment words), and user's interactions with other users (e.g., retweet, reply, friend).

3 News-User Compatibility

3.1 User's compatibility

Figure 1 shows the framework for determining the compatibility between a news article and a Twitter user.

As explained in Section 1, determining a user's compatibility with a news article is currently based on two factors – user's interest match and user's sentiment match. The former is achieved through a simple form of topic mining based on keyword extraction, and the latter is achieved through a rule-based sentiment analysis. Both approaches have been chosen empirically after trying some alternative approaches. Given the two inputs, we first check whether the user matches the news in his/her interest as reflected in his/her tweet texts,

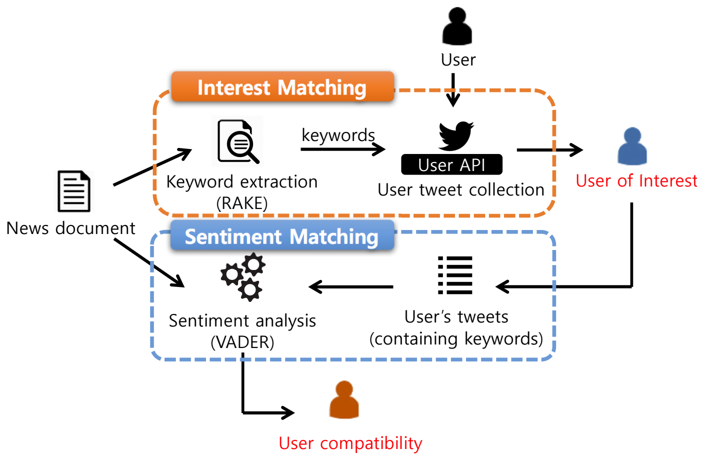

Fig. 1. News-user compatibility framework.

and then check if the user matches the news in the sentiment analysis result as well. Both matches should be confirmed before we determine that the user is compatible with the news. Let us discuss each step in the remainder of this section.

3.2 User's interest match

The decision on whether a user's interest matches the topic in a news article is made according to the following definition.

Definition 1 (Interest match). *Given an article A and a user U, we say U's interest matches A's topic if and only if any of recent tweet texts posted by U contains any of the topic keywords extracted from A.*

In this definition, the cutoff for recency is application-dependent, and all tweets posted by the user within the past recent period are considered for topic keyword match. We use a Boolean keyword match in this work.

As mentioned in Section 1, our definition of interest match is deliberately inclusive by making the match condition existential, and this strategy is necessary in order to overcome the well-known sparsity problem of tweet texts and flimsy tweeting behavior of users. According to Twitter statistics by Sysomos [20], at least 20% of tweets are retweets, and, in our observation, at least 40% of tweet messages contain links. As a result, there is not enough meaningful text in a lot of individual tweet messages. Hashtags, which often are not real words, make the problem worse, as illustrated in Figure 2.

One key issue in this user interest matching is *keyword extraction* from the input article. We first looked into a model-based approach. A topic model was generated by Hierarchical Dirichlet Process (HDP) [22], which is used popularly in topic mining and, given a news article, the topic model was looked up to find matching topic keywords. The performance, however, was inadequate in both speed and accuracy. The reason was that HDP is a classifier, which needs a collection of documents for training and testing and, therefore, it is highly likely that common words irrelevant to the topic are selected from the model. Indeed,

This tweet contains many hashtags that are irrelevant to the topic of the message and there is very little meaningful text content.

Fig. 2. Tweet sparsity caused by hashtags.

when we analyzed a collection of New York Times articles to find topic keywords, HDP extracted words that did not represent the topic of the article.

Thus, we used a simpler approach that does not need a model. The *Rapid Automatic Keyword Extraction (RAKE)* algorithm by Rose et al. [19] was chosen for that purpose. RAKE analyzes a text by considering not only the frequency of each term but also its position and component (e.g., subject, predicate, object) in a sentence. This approach enables the algorithm to extract more meaningful keywords than a simple term frequency analysis can. In addition, the performance is better than using classifiers like HDP because it uses only a single document. RAKE first splits the input text into sentences and then extracts candidate keywords by removing stop words and phrase delimiters. Second, RAKE computes pairwise co-occurrences of the candidate keywords and calculates their scores as the sum of the scores of the words in the keyword. Third, RAKE merges some of the keywords separated in the first step if they have been separated by interior stop words that should be in the same phrase. Finally, from the resulting set of keywords, RAKE returns the keywords ranking one third top scores. In our implementation, we instruct RAKE to return at least three keywords and take the top three from the keywords returned.

3.3 User's sentiment match

The decision on whether user's sentiment reflected in his/her tweets matches the sentiment reflected in the new article is made according to the following definition.

Definition 2 (Sentiment match). *Given a set K of topic keywords extracted from an article A, and a user U with matching interest (according to Definition 1), we say U's sentiment matches A's sentiment if and only if the results of sentiment analysis on U and A are the same, i.e. either both positive or both negative.*

In this definition, sentiment analysis is performed on selected texts, specifically, on only the sentences from an article A that contain *any* of the topic keywords in K, and on only the tweets texts from a user U's tweets that contain *any* of the topic keywords in K. The reason for excluding sentences/tweets that do not contain any of keywords in K is, evidently, that they have adverse effects on the sentiment scores, and the reason for including sentence/tweets that contain any – not all – of the keywords in K is, as in the interest match, to overcome the sparsity problem, especially for tweet texts.

We used a sentiment scoring algorithm [8], which is based on the *Valence Aware Dictionary for sEntiment Reasoning (VADER)* model developed and validated by Hutto and Gilbert [9]. The VADER model is a sentiment lexicon constructed using a combination of qualitative and quantitative methods. More specifically, it is a set of pairs of a lexical feature (i.e., word) and a score, from various resources such as sentiment word banks (e.g., LIWC, ANEW, GI), microblogs, and sentiment-related acronyms. This lexicon is especially accustomed to microblog texts and enables us to achieve impressive results in analyzing the sentiment of tweets.

Given the VADER model, the sentiment scoring algorithm [8] first calculates the triple (i.e., positive, negative, neutral) sentiment matching scores between each word in the input text document and any element in the lexicon whose feature matches the word and accumulates the scores over the words in the input text document. Then, it returns a sentiment score calculated as follows.

$$\text{Sentiment\_score} = \frac{\text{sum}_{\text{pos}} + \text{sum}_{\text{neg}} + \text{sum}_{\text{neu}}}{\text{sum}_{\text{pos}} + |\text{sum}_{\text{neg}}| + \text{sum}_{\text{neu}}} \tag{1}$$

where sum_{pos}, sum_{neg}, and sum_{neu} are the accumulated triple sentiment-matching scores. This is a compound score indicating the overall sentiment of the entire document.

To convert the compound sentiment score to a binary sentiment polarity (i.e., either positive or negative), we introduce a threshold, δ_{neu}, and then assign the polarity to either positive or negative only if the magnitude of the score exceeds the threshold and to neutral otherwise.

As the final step, the sentiment polarities of the news document and the user (with matching interest) are compared and they are determined to be compatible if and only if both have the same polarity. In other words, they are determined to be incompatible in case their sentiment polarities are opposite or at least one has no polarity (i.e., neutral).

Given this sentiment match framework, we have tried two approaches to selecting the keywords that are input to the VADER model-based sentiment scoring algorithm [9]. One is to use the *RAKE-provided keywords* as they are, and one is to process them to identify fewer keywords that are likely to belong to tweets that have the same sentiment polarity. The underlying observation is that two or more keywords co-occurring in news article headlines tend to find *consistent* user's sentiment polarity as a result of the sentiment match according to Definition 2. We thus say that the latter approach uses *co-occurrence-based* keywords. Specifically, we consider different subsets of the three keywords returned from the interest match step and bootstrap New York Times headlines that contain any co-occurring keywords in the different subsets, and then select one or more keywords that occur frequently enough (i.e., term frequency above 80%). In case no such keyword is found, then we use the RAKE-provided keywords.

4 Evaluation

The main objective of the evaluation is to measure the accuracy of our algorithm in determining the compatibility. For this purpose, we conducted gold standard tests using ground truth articles and users. In this section, we compare the accuracies with those obtained when considering only one of the two factors – interest-match-only and sentiment-match-only – in order to validate that we need both. While doing so, we also compare between the two approaches to selecting inputs to the VADER model-based sentiment scoring algorithm (see Section 3.3). Additionally, we provide some examples of output from our algorithm.

All experiments were performed on a Red Hat 4.4.7-1 Linux server with Intel Xeon CPU E5-2667 v2 @3.30GHz and 2GB RAM. Our algorithms were implemented in Python 3.5.2 programming language.

4.1 Experiment setup

Table 1. Ground truth for compatibility accuracy evaluation. URL's of the articles are available at https://github.com/paper-data/user_compatible/blob/master/User_Compatible.pdf.

Seed topic	Sentiment	Article	Compatible users
abortion	negative	US News Instant Answer Youthvoices	LGBT4LifeIRL, DidiJeremie, MichaelKellyIC, ColleenBarry1, misfeet
	postive	Shenvi Salon Media Group NY times	rtraister, HuffPostWomen, ihiccupalot, DonnaHowardTX, KathySchiffer
immigration	negative	Debate opinions Business insider Debate opinions	mkolken, MigrantVoiceUK, reformny, prioritydate, MaddieAndMichi
	postive	American progress The HuffingtonPost Debate opinions	SachaWoolLegal, K_Sreeharsha, jpsimmigrate, DetentionForum, RepGutierrez
Brexit	negative	Netivist The Guardian NY times	eyejosh, Hogmeisster, NYtitanic1999, Australiaunwra6, judithmknott
	postive	The Huffington Post International Business Times Al Jazeera Media Network	moboboandyking, JakubKrupa, massimousai, CarolHope01, bevilwooding

Ground truth Table 1 describes the ground truth. We chose three seed topics abortion, immigration, and Brexit, as they have been the subjects of recent controversy on news media. For each topic, three news articles showing negative-sentiment and three articles showing positive-sentiment were collected from the Internet. News articles whose sentiment scores in absolute value are higher than 0.8 have been selected to assure definite polarity so that the compatibility will be determined by the sentiment scores on the side of the *users* of matching interest.

Additionally, for each topic, we assigned ten users such that five of them are compatible with the positive-sentiment news article and five of them are compatible with the negative-sentiment news article; the compatibility was manually vetted independently of our algorithm. In order to find the users, for each topic, first users whose tweets contain the topic keyword were identified through the

Twitter Advanced Search engine and, then, for each user identified, the user's tweets in the past three months were retrieved through the Twitter User API.

Thus, for each topic, there are exactly five compatible users and 25 incompatible users – specifically, five users have matching interest but non-matching sentiment and 20 users have non-matching interest. These 30 users make a ground truth of an adequate size given the trending nature of tweets reflecting transient interests and sentiments, and they reflect the actual number of users whose tweet messages do not show the sparsity or flimsiness problem mentioned earlier.

As mentioned in Section 3.2, the RAKE algorithm returns at least three keywords in the one-third top scores and we selected the top three keywords from them. When we did it for each news article in the ground truth, the keywords selected for all news articles altogether covered all of the three topics used in the ground truth.

Parameter The threshold parameter δ_{neu}, used in sentiment determination (Section 3.3), is the minimum sentiment score required to assign a polarity to the sentiment score. Its value was set to 0.4 for our algorithm with the co-occurrence-based keyword selection and 0.6 for our algorithm without it and to 0.4 for the sentiment-only algorithm. (Note that the interest-match-only algorithm does not need this parameter.) These values were determined as a result of manually tuning the algorithm outputs against the ground truth.

4.2 Experiment results

Compatibility accuracies Table 2 shows accuracies resulting from four algorithms: interest-match-only, sentiment-match-only, both matches with RAKE keywords without filtering (called "both-RAKE"), and both matches with filtered keywords (called "both-filtered"). Our algorithm with keyword filtering outperforms the interest-match-only by 2.07 times in f-score (resulting from 3.00 times in precision and 0.94 times in recall), and outperforms the sentiment-match-only by 2.93 times in f-score (resulting from 3.55 times in precision and 1.77 times in recall). Table 2 also shows that filtering the keywords from RAKE in our algorithm improves the f-score by 1.29 times over using the RAKE keywords as they are.

We see that the precision increases from interest-match-only or sentiment-match-only to both-RAKE and further increases to both-filtered, which is consistent with the way f-score increases across these algorithms. For recall, interest-match-only shows 1.0 for all news articles. This is from that every compatible user matches in the interest as one of the required conditions. In both-RAKE and both-filtered, however, the recall is lower than 1.0 because some of the keywords for a user may have different sentiment polarities. Evidently, both-filtered achieves higher recall than both-RAKE because there are fewer such keywords as a result of the filtering.

Compatibility example cases Figure 3 shows example cases of news-user compatibility determination. The text of a news article is shown on the left side and an exemplary tweet posted by each of the three tweet users LGBT4Life, Donna Howard, and Larry Hawk are shown on the right side. (We are showing only one exemplary tweet per user due to space limit.) Two topic keywords

Table 2. Compatibility accuracy evaluation results.

In the interest-match-only column, each accuracy number is with regard to the compatibility between the news article in the same row of the ground truth table (see Table 1) and the ten users assigned to the article's topic. The sentiment-match-only algorithm is not relevant to the topic, and therefore each accuracy number reflects only the five users' tweet sentiment. Thus, in the sentiment-match-only column, each accuracy number is with regard to the five users in the same topic & sentiment row of the ground truth table. Our compatibility algorithms (i.e., both-RAKE, both-filtered) consider both interest match and sentiment match. Each accuracy number is with regard to the compatibility between a news article and the five users assigned to the article's topic & sentiment in the same row of the ground truth table.

Seed topic	Sentiment	Interest match only			Sentiment match only			both-RAKE			both-fltered		
		Preci.	Rec.	F-Scr.	Preci.	Rec.	F-Scr.	Preci.	Rec.	F-Scr.	Preci.	Rec.	F-Scr.
abortion	negative	0.22	1.00	0.36	0.27	0.33	0.27	0.60	0.43	0.50	1.00	1.00	1.00
		0.31	1.00	0.48				0.40	0.80	0.53	1.00	1.00	1.00
		0.29	1.00	0.45				0.60	0.60	0.60	1.00	1.00	1.00
	positive	0.21	1.00	0.34	0.32	1.00	0.48	0.24	0.83	0.37	1.00	0.83	0.91
		0.21	1.00	0.34				0.29	0.83	0.43	1.00	0.83	0.91
		0.35	1.00	0.52				0.60	1.00	0.75	1.00	0.83	0.91
immigration	negative	0.26	1.00	0.42	0.12	0.20	0.15	0.83	1.00	0.91	0.71	1.00	0.83
		0.25	1.00	0.40				0.57	0.80	0.67	0.50	0.80	0.62
		0.17	1.00	0.29				0.36	0.80	0.50	0.50	1.00	0.67
	positive	0.25	1.00	0.40	0.11	0.40	0.17	0.50	0.67	0.86	0.71	0.62	0.82
		0.29	1.00	0.45				0.44	0.80	0.57	0.56	1.00	0.71
		0.21	1.00	0.34				0.50	1.00	0.67	0.71	1.00	0.83
Brexit	negative	0.33	1.00	0.50	0.35	0.60	0.44	0.83	1.00	0.91	0.71	1.00	0.83
		0.33	1.00	0.50				0.71	1.00	0.83	0.71	1.00	0.83
		0.23	1.00	0.37				0.63	1.00	0.77	0.71	1.00	0.83
	positive	0.29	1.00	0.45	0.14	0.67	0.21	0.50	1.00	0.67	0.71	1.00	0.83
		0.23	1.00	0.37				0.38	1.00	0.56	0.71	1.00	0.83
		0.22	1.00	0.36				0.56	1.00	0.71	0.71	1.00	0.83
Arithmetic average		0.26	1.00	0.41	0.22	0.53	0.29	0.53	0.86	0.66	0.78	0.94	0.85

"murder" and "abortion" extracted from the news article are shown in the gray box `A` and are highlighted in the news article and the tweets. The sentiment scores are shown in red color – labeled `B` for the news article and `C-1` and `C-2` for tweet messages.

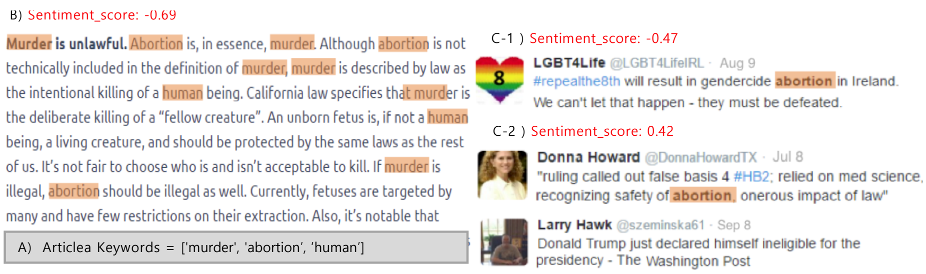

Fig. 3. Compatibility example cases.

The news article's sentiment is negative. First, the user LGBT4Life's tweet contains a matching topic keyword "abortion", so this user matches the news in the interest, and this user's sentiment is negative, which matches the news article's sentiment. Therefore, this user is compatible with the news article. Second, the user Donna Howard's tweet also contains a matching topic keyword "abortion", hence matches in the interest, but this user's sentiment is positive, which is opposite to the news article's sentiment. Therefore, this user is incompatible with the news article. Third, the user Larry Hawk's tweet does contain either topic keyword, and hence this user's interest does not match the news article's topic. Therefore, this user is incompatible with the news article. Our algorithm made correct compatibility decisions in all three cases.

5 Conclusion

The notion of news-user compatibility introduced in this paper enables user profiling, which has many practical applications. The proposed framework combines interest match and sentiment match as two steps for compatibility determination, and each step can be instantiated using different algorithms. In this paper, we used the RAKE algorithm for interest match, to extract topic keywords from the news article for use in matching against the tweets written by the user. Then, we used the VADER model-based sentiment analysis algorithm for sentiment match, to determine the user's sentiment reflected in the user's tweets containing the keywords, where the keywords from RAKE are filtered based on their co-occurrence frequencies in news headlines. Gold standard test results show that considering both interest match and sentiment match and avoiding sentiment ambiguities through co-occurrence-based keyword filtering are instrumental in achieving an average 85% f-score over all ground truth cases.

One further work is to take more factors than interest and sentiment, such as the user's language, geolocation, timeline, and hashtags, into the framework for more accurate compatibility matching. Currently in the plan is the next phase framework, which aims to find users that are compatible with a given news article. Since there is a huge number of users in online social media, iterative processing is needed in order to first find tweets based on semantic and sentimental "cues" and then find users who posted those tweets and then find additional tweets posted by those users, etc. while progressively refining the compatibility scores of the users found thus far.

References

1. Alvarez-Melis, D., Saveski, M.: Topic modeling in twitter: Aggregating tweets by conversations. In: 10th International AAAI Conference on Weblogs and Social Media. pp. 519–522. Cologne (2016)
2. Campbell, W., Baseman, E., Greenfield, K.: Content+context networks for user classification in Twitter. In: Frontiers of Network Analysis: Methods, Models, and Applications Workshop at Neural Information Processing Systems. Lake Tahoe (2013)
3. Cataldi, M., Di Caro, L., Schifanella, C.: Emerging topic detection on Twitter based on temporal and social terms evaluation. In: 16th ACM Conference on Knowledge Discovery and Data Mining. pp. 4:1–4:10. Washington D.C. (2010)
4. Chen, F., Mirisaee, S.H.: Do topic-dependent models improve microblog sentiment estimation? In: 8th International AAAI Conference on Weblogs and Social Media. Ann Arbor (2014)
5. Dan, O., Feng, J., Davison, B.D.: A bootstrapping approach to identifying relevant tweets for socialTV. In: 5th International AAAI Conference on Weblogs and Social

Media. Barcelona (2011)
6. Gutierrez, F.J., Poblete, B.: Sentiment-based user profiles in microblogging platforms. In: 26th ACM Conference on Hypertext and Social Media. pp. 23–32. Guzelyurt (2015)
7. Hong, L., Davison, B.D.: Empirical study of topic modeling in Twitter. In: Workshop on Social Media Analytics. pp. 80–88. Washington D.C. (2010)
8. Hutto, C., Gilbert, E.: VADER-Sentiment-Analysis. https://github.com/cjhutto/vaderSentiment (2014)
9. Hutto, C.J., Gilbert, E.: vader: a parsimonious rule-based model for sentiment analysis of social media text. In: 8th International AAAI Conference on Weblogs and Social Media. Ann Arbor (2014)
10. Kapanipathi, P., Jain, P., Venkataramani, C., Sheth, A.: User interests identification on Twitter using a hierarchical knowledge base. In: 11th European Semantic Web Conference. pp. 99–113. Anissaras (2014)
11. Lim, K.W., Buntine, W.: Twitter opinion topic model: Extracting product opinions from tweets by leveraging hashtags and sentiment lexicon. In: 23th ACM Conference on Information and Knowledge Management. pp. 1319–1328. Shanghai (2014)
12. Lim, K.H., Datta, A.: Interest classification of twitter users using wikipedia. In: 9th International Symposium on Open Collaboration. pp. 22:1–22:2. Hong Kong (2013)
13. Liu, B.: Sentiment Analysis and Opinion Mining. Morgan & Claypool Publishers, San Rafael (2012)
14. Michelson, M., Macskassy, S.A.: Discovering users' topics of interest on Twitter: a first look. In: 4th workshop on Analytics for noisy unstructured text data. pp. 73–80. Toronto (2010)
15. Naveed, N., Gottron, T., Kunegis, J., Alhadi, A.C.: Searching microblogs: Coping with sparsity and document quality. In: 20th ACM Conference on Information and Knowledge Management. pp. 183–188. Glasgow (2011)
16. Pak, A., Paroubek, P.: Twitter as a corpus for sentiment analysis and opinion mining. In: 7th international conference on Language Resources and Evaluation. pp. 17–23. Valletta (2010)
17. Pennacchiotti, M., Popescu, A.M.: A machine learning approach to twitter user classification. In: 5th International AAAI Conference on Weblogs and Social Media. Barcelona (2011)
18. Qadir, A., Riloff, E.: Learning emotion indicators from tweets: Hashtags, hashtag patterns, and phrases. In: 19th conference on Empirical Methods in Natural Language Processing. pp. 1203–1209. Doha (2014)
19. Rose, S., Engel, D., Cramer, N., Cowley, W.: Automatic keyword extraction from individual document. Text Mining: Applications and Theory pp. 1–20 (2010)
20. Sysomos: Replies and Retweets on Twitter. https://sysomos.com/inside-twitter/twitter-retweet-stats (2010)
21. Tan, C., Lee, L., Tang, J., Jiang, L., Zhou, M., Li, P.: User-level sentiment analysis incorporating social networks. In: 17th ACM Conference on Knowledge Discovery and Data Mining. pp. 1397–1405. San Diego (2011)
22. Teh, Y.W., Jordan, M.I., Beal, M.J., Blei, D.M.: Hierarchical Dirichlet processes. Journal of the American Statistical Association 101, 1566–1581 (2006)
23. Tsur, O., Littman, A., Rappoport, A.: Scalable multi stage clustering of tagged micro-messages. In: 21st International World Wide Web Conference. pp. 621–622. Lyon (2012)
24. Zhao, W.X., Jiang, J., Weng, J., He, J., Lim, E.P., Yan, H., Li, X.: Comparing Twitter and traditional media using topic models. In: 33rd European Conference on Information Retrieval. pp. 338–349. Dublin (2011)