Social Semantic Emotion Analysis for Innovative Multilingual Big Data Analytics Markets

D4.9 Social Context Analysis for Emotion Recognition, Initial Version

Project ref. no	H2020 644632
Project acronym	MixedEmotions
Start date of project (dur.)	01 April 2015 (24 Months)
Document due Date	31 December 2015
Responsible for deliverable	UPM
Reply to	J. Fernando Sánchez
Document status	Final

Project reference no.	**H2020 644632**
Project working name	MixedEmotions
Project full name	Social Semantic Emotion Analysis for Innovative Multilingual Big Data Analytics Markets
Document name	MixedEmotions_D4.9_22_12_15_SocialContextAnalysisForEmotionRecognition_UPM
Security (distribution level)	PU
Contractual delivery date	
Deliverable number	D4.9
Deliverable name	Social Context Analysis for Emotion Recognition
Type	Report/Demonstrator/Other
Version	Final
WP / Task responsible	Universidad Politécnica de Madrid (UPM)
Contributors	Fernando Sánchez, Carlos A. Iglesias, Pablo Álvarez
EC Project Officer	Susan Fraser
Document Location	https://confluence.deri.ie:8443/display/mixedem/MixedEmotions+Deliverables+M1-M12

Table of Contents

Executive Summary

This document describes context of intereactions in social networks and its analysis, with an emphasis on its use for Emotion Recognition tasks. We present the architecture of a social context module for this project and its implementation. This implementation includes state-of-the-art metrics of social context aspects, such as influence, which are detailed in the document. We also describe a preliminary experiment with real data captured from Twitter that illustrates the extraction of temporal patterns of emotions in social media. Lastly, we touch on the main challenges to solved in the upcoming months, which will lead to the final version of this deliverable (D4.10).

1 Introduction

Information shared in social networks is not isolated. The meaning of a particular piece of content (e.g. a Tweet, a Facebook status or a blog post) may only be understood when its social context is taken into consideration. In fact, social context has an effect on the behaviour of users in social networks. The behavioral change of a person because of the perceived relationship with other people, organizations and society in general is known as social influence [Sun & Tang 2011]. The strength of social influence depends on many factors such as the strength of relationships between people in the networks, the users of the network or temporal effects. Many applications are based on the social influence phenomenon [Sun & Tang, 2011], such as marketing, advertisement or expert recommendation. Some applications, like viral marketing, aims at influence maximisation.

In contemporary research and industrial practice, most emotion and sentiment analysis systems try to detect collective emotion of social networks through the analysis of user or content attributes. Nevertheless, they usually do not consider how social influence affects users emotions and sentiments. In existing literature, several researchers have studied how social influence can complement sentiment and emotion analysis. West et al. showed that the assumption of homophily in networks can improve polarity detection in short texts [West et al. 2014]. One interesting aspect of the mix of emotion recognition and social network analysis is that the inherent complexity of emotions could also enable us to create richer models of social behaviour. For instance, by mapping traditional categories to a Valence, Arousal and Dominance model [Guerini et al. 2015] we are able to motivate a theory to explain why users in social news have a predilection for certain types of content.

In the MixedEmotions project, we are interested in the analysis and characterisation of the social influence phenomenon for its application in several use cases. First, the detection of most relevant shared contents (e.g. tweets or posts) and users (e.g. influencers) provides a path for micro-analysis of opinions in brand monitoring and content recommendation scenarios. Second, emotion propagation patterns can be used for both analysis as well as prediction of expected social influence of a message. Finally, social features can improve sentiment analysis and emotion techniques. This can be specially relevant in microblogging based social networks such as Twitter, where the short length of the content makes the task very complex.

The rest of the document is structured as follows. First, Section 2 presents a social context model to characterise a social graph and its features. The architecture of a prototype that has been developed to provide the social context of the network Twitter is detailed in Section 3. Our experimentation results, described in Section 4, show the detection of sentiment propagation patterns. Finally, we conclude and describe future work in Section 5.

2 Social context model and features

Social Networks are graphs whose nodes are people and the edges are some kind of relationship between individuals, e.g. a network of cellphone users and their calls. Despite this loose definition, this deliverable will focus on users of online social networks and their interactions.

More specifically, we will use Twitter as the reference social network. All these concepts can be extrapolated to any other microblogging network, and to Facebook, although some restrictions apply. However, other examples tend to suffer of lack of widespread adoption or of a more restrictive interfaces and policies.

There are two main components in social context: users and content. The following subsections present a series of features of both users and content, which we analyse separately. Note that, technically, any information from the social network that is not present in the bare textual content could be considered part of its social context. However, these metrics tend to include richer aspects from the social network, exploiting the graph of relationships and interactions between users and content. Some of these aspects, the more general ones, are already provided by the social network site through its API, such as the number of mentions, favourites or replies. More specific or intensive metrics need to be computed by third parties. This approach is more flexible and powerful, but is obviously limited by the request rate and access restrictions imposed by the social network site.

2.1 User

A user is any entity capable of interacting with other entities in a social network. User information can be split in three groups: profile information, behavior indicators and information about the user network.

2.1.1 Profile

- Biography. The contents of the biography of a user can provide very useful when searching for keywords. It allows us to detect or filter effectively certain topic related users, increasing the likelihood that these users are more influential or popular in those topics.
- Account age. It is more likely that an older account has a high level of influence and popularity, as it has been able to improve and develop its influence and popularity over time, establishing new relationships and increasing its network. Moreover, recent accounts may also be popular, but these tend to be isolated cases. It is a secondary metric and rarely used.
- Number of followers. One of the most common metrics is the number of followers that a certain user has. This is a good metric to estimate the popularity of the user and the network of users who will receive their messages. On the other side, a high number of followers does not always imply high influence, so we may find users with more followers than others but with less influence.
- Number of friends. In practice, this metric is used less frequently than the number of followers. it does not provide much information by itself, as it does not require any kind of acceptance or reciprocity by the other party. However, by combining this metric with the number of followers, we can know how well interconnected a user is with his environment.

- Associated groups. Social media like Facebook allow the creation of user groups, and users may choose to join one or more groups. In other media Twitter, this may be a list of user lists the user has been added to. The difference between both cases would be that in the former groups are global and users choose their own groups, whereas in the latter each user has their own lists and chooses who to add to them. In other words, one is an active selection and the other is a passive one.

2.1.2 Behaviour

- Recent activity. In an environment where messages have a very short lifespan, it is very important to be active and publish often. Knowing the recent activity of a user allows us to filter active accounts and also adds value to users with a constant activity rate. Moreover, this measure loses its value in isolation, because, as in the case of the number of tweets, a very active user does not have to be an influential one.
- Number of posts. The number of posts of a user gives us an idea of the level of activity of this user over time. A large volume of posts makes it more likely that a user stands out, but it is not always the case. On the other hand, this metric is not indicative of popularity or influence by itself. Spammers are an example of users with a large number of posts with little to no impact on the network. For this reason, this feature is usually combined with those of content, in order to give some weighting factor that compensates for these effects.
- Number of mentions. The number of mentions that a user has is usually considered to be a good indicator of popularity and influence, since it involves interaction with the rest of the users. Users with a large number of mentions tend to be very popular and influential in their network. On the other hand, a user who does not have a large number of mentions is not necessary a low influential one, which means this metric only adds positive value to the influence of a user.
- Number of shares. The number of shares a user receives is a good indication of his popularity and influence. A high number of shares implies that the followers of the user usually find interesting or important the content published by the user, which makes him more influential in the network. Besides indicating a high degree of influence, this metric also has a high correlation with other metrics such as the number of followers, the number of mentions or generated traffic.

2.1.3 Network

Several attributes of the social network of a user can be used as an indicator of their influence or impact: betweenness, connectedness, in/out-degree, etc. The network can be constructed in different ways as well, using either user-to-user or user-to-content interactions. Two examples would be using follower/friend relationships and using the comments network. It is common to combine the metrics above with metrics of the content generated by users and their closest network.

2.2 Content

Social media are interesting due to the content being shared between users, which forms the majority of the interactions. There are different types of content (e.g. tweets or status in facebook). Instead of focusing on a particular type, this section provides a high-level description.

2.2.1 Attributes

Content attributes are data associated with the content being shared. Some of the most relevant attributes are:

- Creator. The same content shared by different people has a very different impact on the network. Oftentimes, the audience only depends on the original creator, their network of followers. These factors greatly affect user reactions.
- Keywords. Keywords can be used to group different topics together, and to extract concepts and entities from text. Unfortunately, content does not always directly use keywords.
- Emotion/sentiment. The emotion or sentiment expressed in content can affect other people's reactions. And, as we will see later on, it can reveal other interesting aspects of content and help in prediction.
- Creation time. Content is ephemeral, its value or potential rapidly degrades with time. In other cases, such as social news, the content is tied to specific real life events or contexts.
- Topics. As with keywords, topics are very useful to group content. Moreover, it is a crucial part for emotion recognition, which is domain dependent.
- Media. Textual content may also have multimedia attached. For instance, a Tweet with a picture.
- External links.

2.2.2 Propagation

This kind of information characterises how content is shared and received by other users. Most of these features are dynamic, and we can analyse both the value at a specific point in time or the evolution of this value. Some examples would be: time to achieve a certain amount of reshares, editions to the content over time or evolution of favourites over time.

- Number of reshares. Like in the case of a user, a tweet with a large number of reshares has great popularity and indicates that the users that reshared it find it interesting, making it a good indicator of influence.
- Number of favourites. A high favourite count indicates that a tweet is very popular. On the other hand, a low favourite count does not imply low popularity. In general, people prefer to use a reshares to indicate interest in a tweet, rather than favourite. It is a secondary metric and it can be replaced by the number of reshares in most cases.

- Number of replies. Also known as comments in some networks, this feature measures direct reactions from other users. In contrast to favourites and reshares, which have an inherently positive charge, replies may be negative and include criticism.
- Number of views. In media like Facebook or Youtube, it is possible to determine how many people have accessed the content. This feature is important when combined with others, as it is a proxy measure of neutral reactions.

2.2.3 Network

How content is related to other content and users in a social network often offers valuable insights about the content. The interactions between users and content forms a graph or network that can be studied. This information can be used to detect social media phenomena, such as cascades or memes. These are some of the relationships that can be used, alone or in combination, to form this network: Creation/Modification/Deletion (User → Content); Reshare/Favourite/View (User → Content); and Mention/Reply (Content→Content).

3 Social Context Analysis Module

The Social Context Analysis Module merges the non-contextual annotation of content and the context associated with it under a common tool, allowing a combined analysis.

The module consists of four components as shown in Figure 1: 1) a crawler/scraper that gathers information from social networks and stores it in a graph database; 2) a component that annotates the content stored in the database with emotion, using the textual emotion analysis modules defined in the MixedEmotions platform; 3) a processing component that computes metrics and features from the information in the database; and 4) a REST API to configure the crawler and to get results from the processing component.

The module is Open Source. The source code can be found on Github: http://github.com/mixedemotions/scaner.

3.1 Crawler

The crawler in this module connects to social networks and collects information through their APIs. The crawling parameters can be configured through the API, which will be explained later.

As of this writing, the crawler has access to Twitter through two mechanisms:

- Streaming API, which provides a channel for real-time Twitter events. The main type of event is a Tweet (or Status), of which a retweet is a special case. Tweet events relevant attributes of the content and its creator, such as the user profile or the number of retweets at that time. Other types of event are also provided such as tweet deletions or favourites (of authenticated accounts).
- REST API, for all the information that is not real-time. This includes: user profiles, follower/friend relationships, tweet search, user lists, etc.

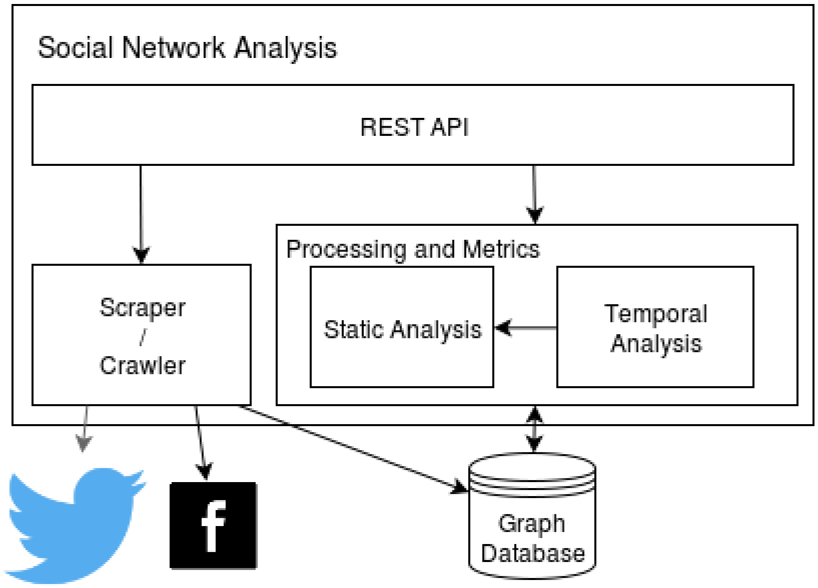

Figure 1: Architecture of the Social Network Analysis module.

3.2 Graph Database

Graph Databases offer a different paradigm from traditional or relational databases and other data stores. In them, data is modelled and queried as graphs, as opposed to the relational database+table approach of traditional databases. Graph Databases are specifically engineered to work with graphs in a very efficient way. This is key when dealing with big datasets, as is the case of MixedEmotions.

In particular, we have chosen to base our implementation on OrientDB. OrientDB offers a mixed model between a pure graph database and a document oriented database such as MongoDB.

3.3 Social Context Analysis

This component of the module provides all the methods to extract the social context attributes described in Section 2, using the data stored in the Graph Database. As of this writing, the following metrics have been implemented to add to the attributes provided by Twitter API, based in previous studies [Noro et al 2013][Noro et al 2014]:

- Tweet Rate Score: This metric measures the proportion of tweets related to the topic that a user posts or retweets. We use the following formula to calculate it:

$$TR(u) = \frac{|t|t \epsilon T \wedge t.user.id = u.id|}{|Total(u))|}$$

Where t is a tweet posted by the user that is relevant to the topic and Total(u) are the total amount of tweets posted by the user u during the topic search duration.

- Follow Relation Factor: This metric shows how well a user is interconnected with the rest of users. In order to calculate it, we create a relation matrix, to build an adjacency matrix.

Then we multiply this matrix to obtain the score. The formulas used are the following ones:

$$A_f(u_i, u_j) = \begin{cases} 1 & \text{if } u_i \text{ follows } u_j \\ 0 & \text{otherwise} \end{cases}$$

$$B_f(u_i, u_j) = \begin{cases} \frac{A_f(u_i,t_j)}{\sum_k A_f(u_i,u_k)}(1-d) + \frac{d}{|U|} & \text{if } \sum_k A_f(u_i, u_k) \neq 0 \\ \frac{1}{|U|} & \text{otherwise} \end{cases}$$

$$f = B_f^T f$$

Where |U| is the total amount of users. At the end, the score is normalized.

$$FR(u_i) = \frac{f(i)}{max_k f(k)}$$

- User Influence Score: This metric measures "how much attention" a user gets from the rest of the users.

$$u = B_t^T t \qquad t = B_a^T u$$

We create three matrices with the relationships between the users and the tweet:

$$A_t(t_i, u_j) = \begin{cases} 1 & \text{if } t_i \text{ is posted/retweeted by } u_j \\ 0 & \text{otherwise} \end{cases}$$

$$A_r(u_j, t_i) = \begin{cases} 1 & \text{if } u_j \text{ retweets/replies to } t_i \\ 0 & \text{otherwise} \end{cases}$$

We transform them in order to get two relationship matrices:

$$B_t(t_i, u_j) = \frac{A_t(t_i, u_j)}{\sum_k A_t(t_i, u_k)}$$

$$B_a(u_j, t_i) = \begin{cases} \frac{A_r(u_j,t_i)}{\sum_k A_r(u_j,t_k)}(1-d) + \frac{A_s(u_j,t_i)}{\sum_k A_s(u_j,t_k)}d & \text{if } \sum_k A_r(u_j, t_k) \neq 0 \\ \frac{A_s(u_j,t_i)}{\sum_k A_s(u_j,t_k)} & \text{otherwise} \end{cases}$$

Now, in order to get the scores, we do the following:

$$u_0 = (\frac{1}{|U|}, \frac{1}{|U|}, ..., \frac{1}{|U|}) \qquad t_0 = (\frac{1}{|T|}, \frac{1}{|T|}, ..., \frac{1}{|T|})$$

$$k = 1$$

$$t_k = B_a^T u_{k-1} \qquad u_k = B_t^T t_k$$

$$k = k + 1$$

The process is repeated until k = 10000
Lastly, the user influence score is normalized.

$$UI(u_j) = \frac{u(j)}{max_k u(k)}$$

- User Relevance Score: This metric is the combined score of the Tweet Rate Score, the User Influence Score and the Follow Relation Factor. We use the previously calculated metrics to obtain the score.

$$UserRel(u) = TR(u)^{w_r} \times UI(u)^{w_i} \times FR(u)^{w_f}$$

Where w_r,w_i and w_j are weights for the different metrics.

- Tweet Influence Score: This metric measures the "amount of attention" that a tweet receives from the users. This metric is calculated along side with the User Influence Score.

$$t = B_a^T u$$

- User Impact: This metrics measures the ability of a user to improve the relevance of a tweet depending on their influence. To calculate it we use the following formula:

$$Impact(u) = \begin{cases} \frac{UI(u)}{|Relate(u)|+\sigma_i} \times (1-d) + \frac{UI(u)}{|T|} \times d & \text{if } |Relate(u)| > 0 \\ \frac{UI(u)}{|T|} & \text{otherwise} \end{cases}$$

  Where *UI(u)* is the user influence of user *u*, *Relate(u)* is the set of tweets that user u has retweeted or replied to, and *T* is the set of all tweets in the search.

- User Voice: This metrics measures the ability of a user to posts or retweets influential tweets.

$$Voice_t(u) = \frac{1}{|Tweet(u)| + \sigma_v} \sum_{t \epsilon Tweet(u)} TI(t)$$

$$Voice_r(u) = \frac{1}{|Retweet(u)| + \sigma_v} \sum_{t \epsilon Retweet(u)} TI(t)$$

  Where *Tweet(u)* is the set of all tweets posted by user *u* and *Retweet(u)* is the set of all tweets that the user *u* has retweeted. σ is an smoothing factor.

- Tweet Relevance Score: This metric shows the relevancy of a tweet based on the "voice" of the original user and the impact of the users that have posted, retweeted or replied to this tweet. We use the following formula:

$$TweetRel(t) = \alpha \times VR(t) + (1 - \alpha) \times IR(t)$$

$$VR(t) = Voice(t.user))$$

$$IR(t) = \sum_{u \epsilon Related(t)} Impact(u)$$

Where α is a weight factor.

- OpenInfluence: This metric show the relevance of a user based on his popularity (number of followers and mentions he receives) and his influence, which is calculated using the number of tweets the user posts per week, the number of retweets that his tweets receives per week and the amount o users that receive or watch his content.

$$r = \log_2(f + m_w) + \log_2((f * t_w) + (\bar{a}r * r_w))$$

Where f are the followers of the user, m_w are the mentions per week of the user, t_w are the tweets posted per week, ar is the retweet audience and r_w are the retweets per week.

3.4 REST API

The REST API is the interface between users and the system. It provides methods to control the crawling parameters and the metrics that are being applied to that data. It also exposes the results of the analysis.

4 Experimentation

In parallel with the definition of the pilots in the MixedEmotions project, we started the definition of set of experiments in order to validate the architecture and social models used in the project.

Our first experiment aims to measure the impact of emotions in the spread of news on social media. We set out to answer the question “does the emotion of a Tweet affect how it propagates?”. It is motivated by recent studies that show that negative and positive information spreads differently [Park et al 2012], and others that have unveiled distinctive patterns of propagation in online media [Yang et al. 2011].

The experiment focuses on social news. Social News is a particular type of social media that is limited to news and events. This content is usually posted by accounts that are associated with a traditional or online news media, and it is linked to an extended version of the article. For instance, the Twitter account of Deutsche Welle (@dwnews) tweets about every new post on their website.

To collect the dataset, we compiled a list of influential social news accounts. The final list includes a total of 40 Twitter accounts such as The Guardian (“@guardian”) and The New York Times (@nytimes). We used the Streaming API to track their activity between 17 March until 20 April 2015. This activity includes tweets, retweets and mentions, which we had to process for the analysis. In total, we collected 329,505 original tweets and 19,383,04 retweets by 756,218 users.

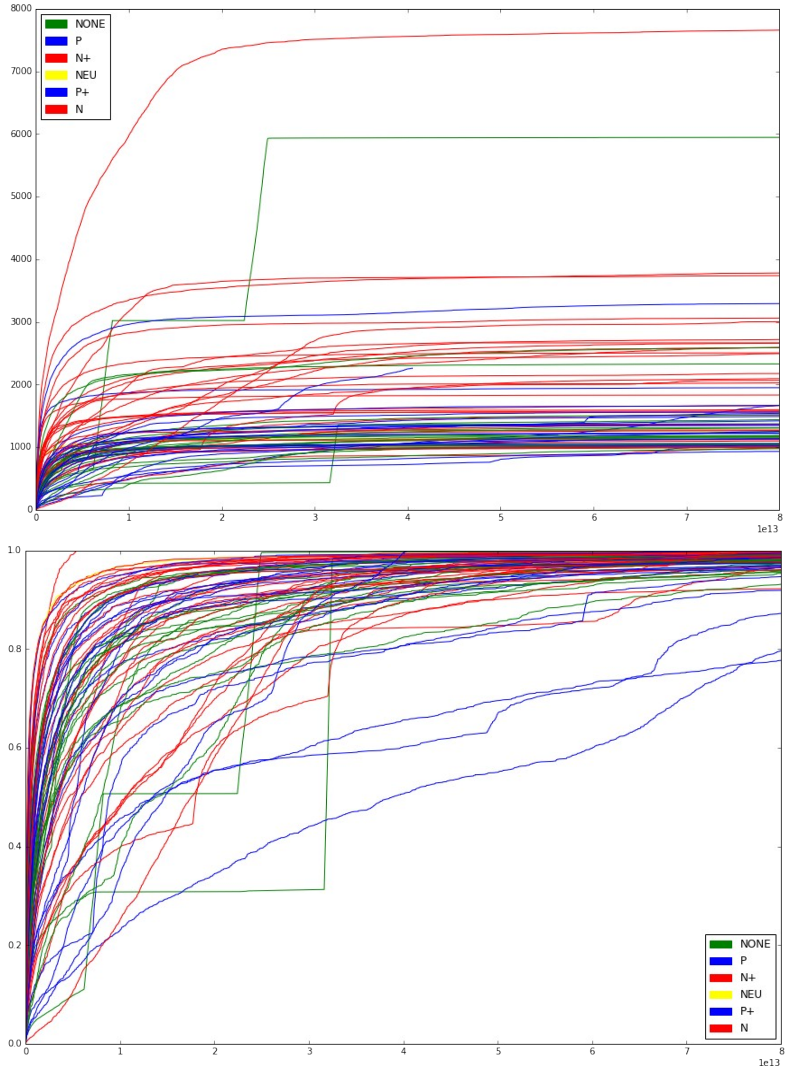

Figure 2: Evolution of number of favourites of the most important tweets since their creation. Below, the same evolution, normalised. Spikes are due to favourite count being sampled when there is a retweet.

we performed sentiment and emotion analysis of the most popular tweets, as measured by the sum of retweets and favourite count.

The core of the experiment consists on comparing the temporal evolution of retweets and favourite count of the selected social news tweets. Thus, the first part was to calculate basic parameters of that temporal evolution to compare different tweets. For every tweet, we got the maximum number of retweets, maximum number of favourites, and the time since the creation of the tweet that it took to reach 25, 50, 75 and 90% of those maximums.

Our preliminary findings are that we have been able to detect sentiment propagation patterns but emotion propagation evolution has not achieved yet conclusive results. This may be due to the fact that the results of the emotion analysis are very poor, which can be observed by manual inspection of the data. This hints that the corpora available for emotion annotation in this domain are not enough to get an analysis subtle enough to detect emotions from social news, and new resources would be needed in order to find patterns for different emotions. Nevertheless, when using polarity and subjectivity detection services, we were able to notice notable differences between positive and negative content, both in audience reached and temporal evolution. On average, it took negative tweets half the time to reach twice the retweet count of the positive ones. The visual representation can be seen This confirms previous research [Rozin et al 2014] and the results obtained in similar studies on Facebook [Kramer et al 2014]. A visual representation of the evolution of the most popular tweets can be seen in the figures below.

Additionally to these tweets, we have gathered follower/friend relationships. Given the size of the network, we selected those users that retweeted 100 or more times, which results in a list of the 925 top retweeters. From those users, we collected their followers, a total of 7,730,067 follower-followed links. This is the network of users that were potentially exposed to the original retweets. The results are currently being used to test the Social Analysis Module with real world data.

5 Conclusions and future work

There is valuable information hidden in the interactions between users in social networks that has the potential to improve emotion recognition in certain scenarios. This information ranges from the simpler features provided by the social network sites themselves, to the more elaborate social network analysis metrics. This is a field that is slowly starting to get attention, mostly in the more constraint domain of polarity detection. MixedEmotions can help foster the use of these hybrid techniques by showing the technology in its pilots. The creation of a module to encompass and facilitate the social network analysis is a step forward in that direction.

However, if we aim to discover the best social features for emotion recognition, we need to refine our current services to get finer results. For that, we need to start making use of the simpler social features, and also create more quality resources for emotion recognition. Our preliminary results with the test corpora support that view. Then, we will be in a position to create better models of user behaviour, which we could use to improve the annotation of their content. This will be the aim of the final version deliverable for this task.

References

- Aral, S., Walker, D., 2011. Creating social contagion through viral product design: A randomized trial of peer influence in networks. Management Science 57, 1623–1639.
- Bakshy, E. et al, 2011. Everyone's an influencer: quantifying influence on twitter, Proceeding
- WSDM '11 Proceedings of the fourth ACM international conference on Web search and data mining.
- Bermingham, A., Conway, M., McInerney, L., O'Hare, N., Smeaton, A.F., 2009. Combining social network analysis and sentiment analysis to explore the potential for

online radicalisation, in: Social Network Analysis and Mining, 2009. ASONAM'09. International Conference on Advances in. IEEE, pp. 231–236.

- Cheng, J., Adamic, L., Dow, P.A., Kleinberg, J.M., Leskovec, J., 2014. Can Cascades Be Predicted?, in: Proceedings of the 23rd International Conference on World Wide Web, WWW '14. ACM, New York, NY, USA, pp. 925–936. doi:10.1145/2566486.2567997
- Coviello, L., Sohn, Y., Kramer, A.D.I., Marlow, C., Franceschetti, M., Christakis, N.A., Fowler, J.H., 2014. Detecting Emotional Contagion in Massive Social Networks. PLoS ONE 9, e90315. doi:10.1371/journal.pone.0090315
- Fan, R., Zhao, J., Chen, Y., Xu, K., 2014. Anger is more influential than joy: Sentiment correlation in Weibo. PloS one 9, e110184.
- Guerini, M., Staiano, J., 2015. Deep Feelings: A Massive Cross-Lingual Study on the Relation between Emotions and Virality. arXiv preprint arXiv:1503.04723.
- Hillman, R; Trier, M, 2013, Influence and Dissemination of Sentiments in Social Network Communicaton Patterns, Proceedings of the 21st European Conference on Information Systems.
- Hirst, M., Harrison, J., Mazepa, P., 2014. Communication and new media: From broadcast to narrowcast. Oxford University Press.
- Jenders, M., Kasneci, G., Naumann, F., 2013. Analyzing and Predicting Viral Tweets, in: Proceedings of the 22Nd International Conference on World Wide Web, WWW '13 Companion. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, pp. 657–664.
- Khabiri, E., Hsu, C.-F., Caverlee, J., 2009. Analyzing and Predicting Community Preference of Socially Generated Metadata: A Case Study on Comments in the Digg Community., in: ICWSM.
- Lee, C.-H., Chien, T.-F., Yang, H.-C., 2011. An automatic topic ranking approach for event detection on microblogging messages, in: 2011 IEEE International Conference on Systems, Man, and Cybernetics (SMC). Presented at the 2011 IEEE International Conference on Systems, Man, and Cybernetics (SMC), pp. 1358–1363. doi:10.1109/ICSMC.2011.6083847
- Lerman, K., Galstyan, A., 2008. Analysis of social voting patterns on digg, in: Proceedings of the First Workshop on Online Social Networks. ACM, pp. 7–12.
- Seth, A., Zhang, J., Cohen, R., 2008. A multi-disciplinary approach for recommending weblog messages, in: Proc. AAAI Workshop on Enhanced Messaging.
- Siersdorfer, S., Minack, E., Deng, F., Hare, J., 2010. Analyzing and predicting sentiment of images on the social web, in: Proceedings of the International Conference on Multimedia. ACM, pp. 715–718.
- Sun, J, Tang, J., 2011. A Survey of Models and Algorithms for Social Influence Analysis, in Social Network Data Analytics, pp 177-214, Springer.
- Yang, H., Willis, A., de Roeck, A., Nuseibeh, B., 2012. A Hybrid Model for Automatic Emotion Recognition in Suicide Notes. Biomed Inform Insights 5, 17–30. doi:10.4137/BII.S8948
- Guerini, M., Staiano, J., 2015. Deep Feelings: A Massive Cross-Lingual Study on the Relation between Emotions and Virality. arXiv preprint arXiv:1503.04723.
- Noro, T., & Tokuda, T. (2014). Effectiveness of Incorporating Follow Relation into Searching for Twitter Users to Follow. In Web Engineering (pp. 420-429). Springer International Publishing.
- Noro, T., Ru, F., Xiao, F., & Tokuda, T. (2013). Twitter user rank using keyword search. Information Modelling and Knowledge Bases XXIV. Frontiers in Artificial Intelligence and Applications, 251, 31-48.