Proceedings of the 1st International Workshop on Social Influence Analysis (SocInf 2015)
July 27th, 2015 - Buenos Aires, Argentina

A Novel Metric for Assessing User Influence based on User Behaviour

Antonela Tommasel and Daniela Godoy
ISISTAN Research Institute, CONICET-UNCPBA, Tandil, Buenos Aires, Argentina

Abstract

People's influence has been the subject of study of several social and humanities disciplines. Lately, the study of user's influence in micro-blogging platforms arises as an important issue. Although social influence or prestige can be defined as the potential or ability of an individual to engage others in a certain act, or to induce others to behave in a particular manner, there is no global consensus on what means to be an influential user. This work aims at shedding some light on how to assess user influence by proposing a novel metric of user influence based on analysing user behaviour regarding both content-based and topological factors. The metric does not only consider each user individually, but also aims at assessing the interactions with his/her neighbourhood. The statistical analysis performed confirmed that only analysing the topological factors is not sufficient for accurately assessing the influence of users. Instead the published content and its influence over the neighbourhood of users has to be also analysed. A comparison with a human assessment of user influence showed that the factors considered by the proposed metric are truly relevant for assessing people's influence.

1 Introduction

Several disciplines, such as sociology, communication, marketing and political sciences have tackle the study of people and their influence [Rogers, 2003; Katz and Lazarsfeld, 2005]. The notion of people's influence plays a crucial role in businesses and in the functioning of societies. For example, it could modify the spread of fashion or voting patterns [Gladwell, 2000; Keller and Berry, 2003]. Furthermore, the study of influence patters could help to understand how trends and innovations are adopted, and to the design of more effective publicity campaigns [Cha *et al.*, 2010]. The rapid growth and exponential usage of social digital media increased the popularity of micro-blogging platforms, which have become an important part of the daily life of millions of users scattered across the world. As a result, the study of users' influence in the context of micro-blogging platforms arises as an important issue.

Social influence or prestige can be defined as the potential or ability of an individual to engage others in a certain act, or to induce others to behave in a particular manner. In micro-blogging sites, the definition traditionally relies on status attributes such as the number of followers (i.e. the size of the influence group of a certain user), the number of re-tweets (i.e. the ability to generate attractive content to be distributed), and the number of mentions (i.e. the ability to engage other users in a conversation). However, having a high number of followers, which would imply a high level of popularity, it is not sufficient for also being influential in terms of triggering social responses as retweets or mentions [Cha *et al.*, 2010]. Moreover, the most influential users are influential over several topics, but such influence is obtain only through a concentrated effort of posting tweets related to only one topic.

In this context, analysing the behaviour of users regarding the diffusion of information can be useful for assessing their influence. Several authors have proposed characterisations of users based on behavioural patterns observed through not only topological features [Java *et al.*, 2007; Krishnamurthy *et al.*, 2008], but also social and content-related features [Tinati *et al.*, 2012]. Such categorisations analyse the number of published posts, the type of posts (original posts, replies in conversations, retweets), the number of times that posts were retweeted, the proportion of retweeted posts, the proportion of followers, and the number of interactions with the neighbourhood, among others. Consequently, the influence of users can be estimated according to certain behavioural patterns. For example, users who share a large number of posts, which are highly retweeted, and also have more followers than followees could be regarded as highly influential users. On the other hand, users who rarely publish posts and have a larger proportion of followees than followers could be regarded as not influential.

There are also several commercial metrics that claim to be able to assess the influence of users, such as *Klout*[1] and

[1] http://klout.com/

Kred[2], among others. However, they have received several critics and have been the focus of several controversies regarding how the measurements are computed or the effect that spam-bots might have on the algorithms. As most of the commercial measures do not publicly state how scores are computed, they are not accessible for scrutiny or reproduction, which might compromise their trustworthiness [Gaffney and Puschmann, 2012].

Considering that there might be no consensus on what means to be an influential user, this work aims at shedding some light on how to assess user influence by proposing a novel metric based on analysing user behaviour regarding the patterns of information diffusion, i.e. it considers both content-based and topological factors. The metric does not only consider each user individually, but also aims at assessing the interactions with his/her neighbourhood. Then, a statistical analysis is performed for comparing the novel metric with traditional means for assessing user influence (such as In Degree or followee/follower ratio), commercial metrics and a human assessment of user influence.

The rest of this paper is organised as follows. Section 2 presents several characterisations of users regarding their role on the diffusion of information. Section 3 presents and defines the proposed metric for estimating the influence of *Twitter* users based on the patterns of user behaviour regarding the information diffusion process. Section 4 describes the analysis carried out using *Twitter* data. Section 5 discusses related research. Finally, Section 6 summarises the conclusions drawn from this study and presents future lines of work.

2 Information-based User Characterisation

Several studies [Java *et al.*, 2007; Krishnamurthy *et al.*, 2008] have characterised users according to their behaviour in the information diffusion process by classifying them into three categories: *Information Sources* or *Broadcasters*, *Friends* or *Acquaintances*, and *Information Seekers*. *Information Source*s are those users who have a greater proportion of followers than followees (i.e. they are followed by more users than they follow) and also publish valuable and relevant content on a regular basis. *Friends* are those users who have a balanced number of followees and followers, without necessarily implying the presence of reciprocal relationships. Finally, *Information Seekers* are those users who rarely publish content and have a greater proportion of followees than followers, aiming at receiving updates. Furthermore, Tinati *et al.* [2012] characterised users into five dimensions (*Idea Starters*, *Amplifiers*, *Curators*, *Commentators* and *Viewers*) according to their social and psychological behaviour, and how this behaviour affects their posting and communication activity on *Twitter*. Each dimension can be associated to the categories identified in [Java *et al.*, 2007; Krishnamurthy *et al.*, 2008]. More importantly, these characterisations can be used for assessing the influence of users:

Idea Starters are those users who are highly engaged with the media. As they tend to start conversations, the majority of their posts correspond to original content and to be highly retweeted suggesting their influence over their followers. These users tend to interact with a limited and selected group of users, ensuring high quality relations with them. They share the characterisation proposed for *Information Sources*.

Amplifiers are those users who share ideas and opinions posted by other users. They tend to have a greater number of followers than followees. They interact with *Idea Starters* and share their ideas with a more visible audience. As a result, the majority of their posts correspond to retweets of *Idea Starters*. Posts are also highly retweeted by their followers. They share the characterisation proposed for *Information Sources*.

Curators are those users who interact with both *Idea Starters* and *Amplifiers* by aggregating their ideas together, and helping to clarify the topic of conversation. They tend to have a balanced number of followees and followers, and to interact with a large number of them. They lie in the border between *Information Sources* and *Friends* as they tend to share a lot of content (as an *Information Source*) and to interact with a large number of users (as a *Friend*). As the number of interactions with other users increases, the *Curator* behaves as a *Friend*, whereas as the number of interactions with other users decreases, the *Curator* behaves as an *Information Source*.

Commentators are those users who also share the ideas and opinions of other users, but without interrupting the flow of the original conversation or immersing in it. They only want to share content and do not desire to be recognised by their posts. The main difference between *Amplifiers* and *Commentators* is the impact that their content has over their social network, measured by the number of retweets received. As the number of retweets increases the user behaves more as an *Amplifier* and less as a *Commentator*. They can be characterised as *Friends*, as they tend to have a balanced number of followees and followers.

Viewers are those users who do not share nor publish posts. They do not engage on conversations or retweet other posts. Instead, they read or consume large amounts of information. They tend to have a larger number of followees than followers. They share the characterisation proposed for *Information Seekers*.

3 Quantitatively Assessing User Influence

Based on the defined characterisations and dimensions of user behaviour, this section proposes novel definitions for quantitatively analysing the behaviour of users in each of the presented dimensions, and then assessing their influence. The definitions of *Commentators* (those users who also share the ideas and opinions of other users, but without interrupting the flow of the original conversation or immersing in it) and *Viewers* (those users who do not share nor publish posts) are omitted as they can be inferred from the scores of the other dimensions. For example, a low score in the *Amplifier* dimension could indicate the presence of a *Commentator*. On the other hand, low scores in both *Idea Starter* and *Amplifier* dimensions could indicate the presence of a *Viewer*. In all cases, the scores are constrained to the interval $[0, 1]$, and corrections

[2] http://kred.com/

are applied in order to avoid undetermined values.

3.1 Idea Starter

Idea Starters are characterised for posting a greater proportion of original content ($Tweets_{ORIGINAL}$) than the other dimensions. Equation 1 not only assesses the proportion of original posts (first part), but also the impact of those posts in the neighbourhood of the user (second part).

$$\frac{|Tweets_{ORIGINAL}\{Tweets_{ORIGINAL}RT >= \mu - \sigma\}|}{|Tweets|} * \frac{\sum Tweets_{ORIGINAL}RT}{|ReTweets|} \tag{1}$$

The first part considers the ratio between the number of original tweets with a number of retweets superior to the inferior limit of the normal distribution of retweets ($|Tweets_{ORIGINAL}\{Tweets_{ORIGINAL}RT >= \mu - \sigma\}|$, where $Tweets_{ORIGINAL}RT$ is the number of retweets that $Tweets_{ORIGINAL}$ has received, and μ and σ represent the arithmetic mean and standard deviation of the retweet distribution, respectively) and the total number of published tweets ($|Tweets|$). The restriction imposed on the number of retweets assesses whether the received retweets are uniformly distributed over all published tweets or over a small proportion of them. The second part assesses the impact that posts have on the neighbourhood of the user, which is measured as the ratio between the retweets that the original content received ($\sum Tweets_{ORIGINAL}RT$) and the total number of retweeted tweets ($|ReTweets|$). The higher the score, the more the user behaves as an *Idea Starter*, and thus as an *Information Source*.

3.2 Amplifier

They are characterised for posting a greater proportion of retweeted content ($Tweets_{RT}$), engage on conversations ($Tweets_{REPLY}$) or even start conversations by mentioning other users ($Tweets_{MENTION}$). Equation 2 assesses not only the interaction between a user and his/her social network (first part), but also the impact of those posts in such network (second part).

$$\frac{|Tweets_{RT}|+|Tweets_{REPLY}|+|Tweets_{MENTION}|}{|Tweets|} * \frac{\sum Tweets_{RT}RT+\sum Tweets_{REPLY}RT+\sum Tweets_{MENTION}RT}{|ReTweets|} \tag{2}$$

The first part considers the ratio between the added number of retweeted content, conversations and tweets containing mentions, and the total number of published tweets. The second part assesses the impact that posts considered in the first part have on the social network of the user, which is measured as the ratio between the retweets received by the retweeted content ($\sum Tweets_{RT}RT$), the conversations ($\sum Tweets_{REPLY}RT$) and the mentions of other users ($\sum Tweets_{MENTIONS}RT$), and the total number of retweets. Also, the second part aids in the accurate differentiation between *Amplifiers* and *Commentators*. A high score in the second part might indicate the presence of an *Amplifier*, whereas a low score might indicate the presence of a *Commentator*. The higher the score, the more a user behaves as an *Amplifier*, and thus as an *Information Source*.

3.3 Curator

They are characterised for interacting with a greater number of users than those characterised by the other dimensions. By default any user can interact with any other user, regardless whether they are actually followers of that other user. Equation 3 assesses not only the number of interactions with other users (first part), but also to what extend a user interacts only with his/her neighbourhood.

$$\frac{|Interactions \in \{Followers \cup Followees\}|}{|Interactions|} * \frac{|Interactions \in \{Followers \cup Followees\}|}{|Followers|+|Followees|} \tag{3}$$

The first part considers the ratio between those interactions that belong to either the follower or followee list ($Interactions \in \{Followers \cup Followees\}$), and the total number of interactions ($Interactions$). Then, the second part considers the proportion of users with whom a certain user interacts regarding the size of the neighbourhood. The higher the number of interactions, the higher the score, and thus the less the user behaves as an *Information Source*.

3.4 Follower/Followee Ratio

In addition to the content-related dimensions, a topological factor can also be considered. The content-related dimensions do not consider the size of the neighbourhood of a user. As a result, two users might achieve the same score but have a totally different neighbourhood. In other words, it is more important a user with a high content-related score and a greater neighbourhood engaged in his/her content (i.e. a greater number of followers) than a user with a high content-related score but a smaller neighbourhood engaged in his/her content. Consequently, the Follower/Followee$_{Ratio}$ (FF_{Ratio}), is proposed to leverage the importance of the neighbourhood size (Equation 4).

$$\frac{|Followers|}{|Followers| + |Followees|} \tag{4}$$

3.5 Information Source Index

Based on the previous metrics, the *Information Source Index* (IS) is defined for numerically characterising users according to their behaviour. The metric denotes to what extent a user can be considered an *Information Source* or an *Information Seeker*. High values of IS denote users behaving as *Information Sources*, whereas low values of IS denote users behaving as *Information Seekers*. For computing the IS index the *Idea Starter*, *Amplifier* and $1 - Curator$ are assigned equal weight and thus combined by means of the arithmetic mean (μ_{IDAC}), as shown in Equation 5.

$$\mu_{IDAC} = \frac{\text{Idea-Starter} + Amplifier + (1 - Curator)}{3} \tag{5}$$

Then, the combination of the content-related dimensions (i.e. μ_{IDAC}) and the topological factor FF_{Ratio} are combined by means of the Harmonic mean for defining the IS, as shown in Equation 6. As the content-based dimensions and the topology factor represent different aspects of user behaviour, they are different kind of elements, and thus cannot be combined by means of the arithmetic mean. Consequently, the Harmonic mean is more adequate for computing the final score. Furthermore, the Harmonic mean is less biased to the presence of small numbers or outliers.

$$IS\left(u_j\right) = \frac{2 * \mu_{IDAC} * FF_{Ratio}}{\mu_{IDAC} + FF_{Ratio}} \quad (6)$$

As *Information Sources* represent those users who are highly engaged with the media, publish valuable and relevant content on a regular basis they could be considered influential users. Furthermore, they tend to engage a great audience of *Amplifiers* and *Commentators* who share and enrich their posts. Due to its relevance, their published content tends to be highly retweeted, which also implies a high number of interactions with their neighbourhood. Additionally, they tend to be highly followed by *Viewers*. As a result, *Information Sources* meet all the requirements for being regarded as influential users. In this context, the influence of users could be measured by means of the IS score. The higher the IS of a user, the higher the influential such user is supposed to be.

4 Data Analysis

This section presents the experimental evaluation performed to assess the effectiveness of the proposed metric. Section 4.1 presents other scores for measuring the influence of users in the context of social networks to which the presented metric was compared. Section 4.2 describes the data collections and data analysis settings regarding the human assessment of user influence. Finally, Section 4.3 presents the results of the data analysis performed.

4.1 Metrics used for Comparison

In order to quantitatively assess the effectiveness of the proposed metric for measuring the influence of users, it was compare to the scores of other related metrics. The first two are commercial metrics, whereas the last two are metrics that can be easily computed with simple data obtained from the users' profiles.

Klout was launched in 2008. It provides a measurement of the online influence of users by combining information extracted from *Twitter*, *Facebook*, *LinkedIn*, *Instagram*, *Google+*, *Flickr*, *Blogger* and *Foursquare*, among others. It is based on three fundamental principles: True Reach (how many users a certain user influences), Amplification (how much a certain user influence other users), and Network Impact (the influence of networks of users).

Kred was launched in 2012. It measures social influence and outreach in *Twitter* and *Facebook*, aiming at assessing the trust and generosity of users. Social influence is measured by assessing the retweets, replies, mentions and new followees a user has, i.e a user receives social influence points every time people interacts with his or her content.

In Degree. Computes the influence of a user as the number of his/her followers. This metric is currently used by many third-party services, such as *TwitterHolic*[3].

Follower/Followee Ratio. Computes the influence of a user as the ratio between their followers and followees. A high score indicates that the user has a higher proportion of followers than followees.

[3] http://twitaholic.com/

4.2 Data Analysis Settings

Several data collections were created by manually selecting *Twitter* users who were considered influential or even popular according to the criteria presented by *Twitter Counter*[4] and *WeFollow*[5]. Selected users were grouped according to their topic of influence. In total seven datasets were created: *Argentina* (13 users), *Miscellaneous* (8 users), *Music* (8 users), *Politics* (7 users), *Sports* (7 users), *Technology* (7 users) and *Tv-Movies* (10 users). For every user, all tweets, followees, followers, favourite tweets and user account information were retrieved. All the information was obtained by means of requests to the *TwitterAPI*[6]. Additionally, for each user his or her *Kred* and *Klout score* was obtained during March 2015 by means of the *Kred* [7] and *Klout* [8] APIs respectively.

Considering that there is no consensus on what means to be an influential user [del Campo-Ávila *et al.*, 2013; Gaffney and Puschmann, 2012], this work aims at analysing the human perception of influence. The presented metric was compared not only to several commercial metrics, but also to a human assessment of user influence. Undergraduate and graduate students from an Artificial Intelligence course at UNICEN University (Argentina) were asked to rank the sets of users previously presented according to the perceived influence. The ranking task was materialised by means of a web site [9] in which the students were able to access a brief summary of the profiles and latest tweets posted by each *Twitter* user. All rankings were performed during April 2015. In total, 31 students ranked the users according to their perceived influence. In order to combine the rankings provided by the students, *Twitter* users were assigned the mode of the provided rankings.

Once all users were ranked according to their influence score in each metric, it is possible to quantify how the rank of users varies across the different metrics. The correlation between the different rankings was analysed by means of the Kendall coefficient [Kendall, 1938], which is a statistic used for measuring the association between two measured quantities in the form of lists or rankings. The correlation takes a value between -1 and 1 so that the higher the score, the higher the agreement between the two rankings. A score equal to 0 indicates that the two rankings are independent. The correlation is analysed for the total number of *Twitter* users in each dataset.

4.3 Analysis of Influence Metrics

Figure 1 shows the correlation among the different analysed metrics. As it can be observed, in all cases the correlation between rankings depends on the dataset under consideration. This could indicate that the topic that users publish about might be an important factor to consider in the analysis. As regard the commercial metrics (Figures 1b and 1c), their highest correlations values were found for the *Technology* dataset, where the correlation between *KloutScore* and

[4] http://twittercounter.com/
[5] http://wefollow.com
[6] https://api.twitter.com/
[7] https://developer.peoplebrowsr.com/kred/
[8] https://klout.com/s/developers/home/
[9] https://sites.google.com/site/influenciatwitterusers/

KredScore was higher than 0.8. On the contrary, the lowest correlation among the commercial metrics was found for the *Miscellaneous* dataset. However, only for the *Technology* dataset the p-values were lower than 0.05, which indicates that for all the other datasets the null hypothesis cannot be rejected and thus, it can be stated that both commercial metrics are actually independent. These results reinforce the idea that there is no consensus between the commercial metrics on how the influence of users is computed.

For most datasets, both commercial metrics were not correlated with the metric presented in this work. These results could imply that the content-based and topological dimensions the IS calculation is based upon are not regarded in the same manner or with the same importance by the commercial metrics. Moreover, the commercial metrics could be also based on a different set of dimensions or features. The only exception was for the *Politics* dataset in which the highest correlation value corresponded to the IS. Additionally, the degree of correlation between the commercial metrics and the FF-ratio and In Degree also depended on the considered dataset. For example, regarding *KloutScore*, the highest correlation with In Degree was found for the *Miscellaneous* dataset, whereas the highest statistically significant correlation with the FF-ratio was found for the *Technology* dataset. Interestingly, for the *Technology* dataset, the correlation between In Degree and *KloutScore* was lower than for the FF-ratio, which could indicate that in such topic, it is not highly important the actual number of followers, but the proportion of followers regarding the number of followees.

As regards the IS (Figure 1a), for three datasets (*Argentina*, *Miscellaneous* and *Technology*) the highest correlation values corresponded to the human assessment. This could indicate that the human users agree with the criteria considered by the proposed metric for measuring influence. On the contrary, for two datasets (*Sports* and *TV-movies*) the highest correlation values corresponded to the *KloutScore*. Furthermore, for those datasets the correlations to the other metrics were negative. Note that, when analysing the correlation between FF-ratio, the human assessment and every other metric for those datasets, results were similar. As a result, it can be infer that, in those cases, the human assessment of influence was mostly guided by topological factors, such as the number of followees, and not by the content they post or the impact that such content has in the form of retweets. The highest overall correlations were found for the *Politics* dataset, being the most statistically significant the ones with *KloutScore* and the human assessment. Conversely, the correlation with the In Degree was statistically insignificant. These results highlight the importance of not only considering the topological links, but also the published content and its impact.

Regarding the human assessment (Figure 1d), the highest overall correlations were found for the *Miscellaneous* and *Technology* datasets with the FF-ratio, which could imply the preference for users with a higher proportion of followers. However, this could be also caused by human users not being familiar with the *Twitter* users they had to analyse. It is worth mentioning that for the dataset *Argentina*, the highest correlation value was found for the IS. Furthermore, for this dataset the correlations with both topological metrics were negative,

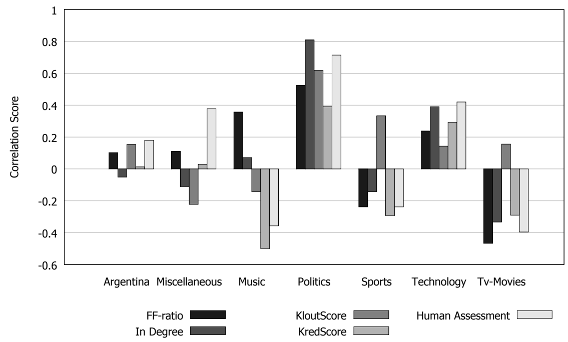

(a) IS index Vs. All

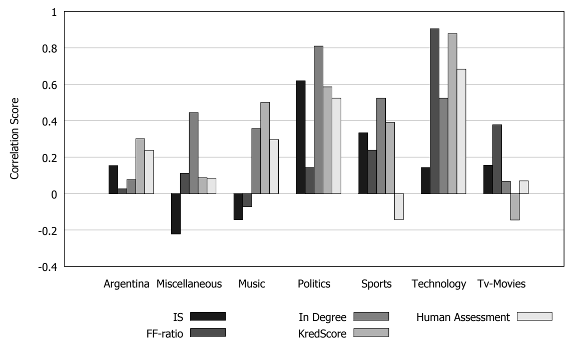

(b) Klout Vs. All

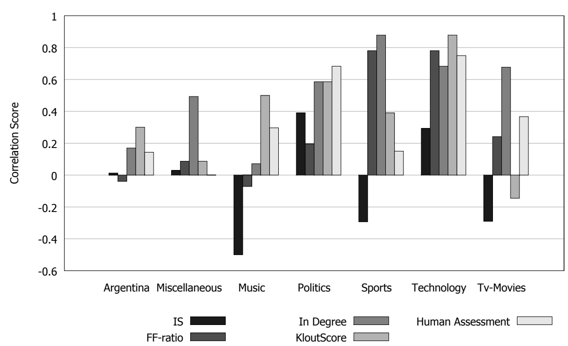

(c) Kred Vs. All

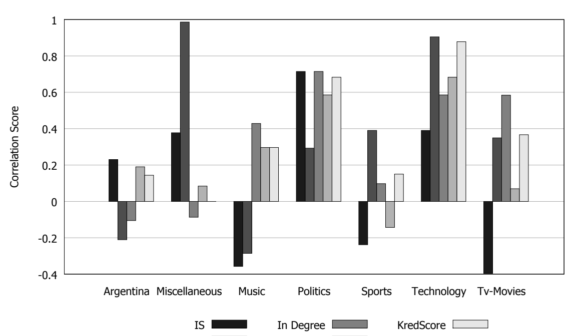

(d) Human Assessment Vs. All

Figure 1: Correlations between metrics across the different datasets

which reinforced the fact that topological factors are not sufficient for assessing the influence of users regarding their real impact or influence in their neighbourhood. These results are of great importance as the human users were highly familiar with all the users in the dataset, and thus these rankings can be regarded as the most accurate ones.

In summary, as in most cases the analysed metrics were not highly correlated, results highlighted the fact that there is no consensus on how the scores are computed and that defining user influence cannot be considered a trivial task. Moreover, in several cases the influence assessment of the different metrics proved to be independent from each other. Furthermore, results seemed to indicate that among the different topics might not be an uniform consensus regarding what means to be influential, which further remarks the fact that there is no unique definition of user influence. Finally, results showed that in specific topics (for example *Music* and *Sports*) the appreciation of human users might be related mostly to the popularity of users measured by means of topological factors.

5 Related Work

Cha *et al.* [2010], compared three measures of user influence: In Degree, Retweet and Mention influence. Results showed a strong correlation between the Retweet and Mention influence. However, In Degree was not strongly correlated to the other two measures, which could imply that the most connected users are not necessarily the ones that are most capable of engaging others in conversations or on spreading their tweets.

Also, Kwak *et al.* [2010] compared three measures of influence: Followers, Retweets and PageRank. Results agreed with those in [Cha *et al.*, 2010]. The three rankings comprised different users, and only 4 users out of 20 appeared in the three rankings. The authors found that the Retweet ranking differed from the other two, which could indicate that not necessarily the most followed users are also the most retweeted ones. Both works highlighted the fact that user influence can be defined from different points of view, which are not necessarily contradictory.

On the other hand, several works [Weng *et al.*, 2010; Yamaguchi *et al.*, 2010] focused on creating more sophisticated approaches for measuring user influence by combining both topological and content-based features. Weng *et al.* [2010] presented *TwitterRank*, an extension of the *PageRank* algorithm, that unlike the original algorithm, considers the topological structure of the network and the topical similarity between users. Experimental results showed that *TwitterRank* was able to improve the algorithm used by *Twitter,* and both the original *PageRank* and Topic-sensitive *PageRank*. Also, the study confirmed the existence of homophily in *Twitter*, justified by the fact that there are users who follow others because they actually have some interest in common and not due to chance.

Alike the previous approach, Yamaguchi *et al.* [2010] presented *TuRank* for measuring users' influence based on both content information and topology. In this case, the content information was considered by analysing how tweets flow among users, i.e. the retweeting phenomenon.Four versions of *TuRank* were compared to 8 ranking schemes, including number of followers and retweets, *PageRank*, HITS [Kleinberg, 1999]. According to the authors, all the other ranking schemes were outperformed by *TuRank* as they only consider topological information, suggesting the importance of considering also content.

As regards commercial metrics, Messias *et al.* [2013] found that they might be vulnerable and easy to manipulate. The authors developed bot accounts, which were able to interact with real users by following them or posting tweets about interesting topics by following different patterns of followee selection and posting activity. Results showed that bot accounts were able to become influential by following simple strategies, reaching similar or higher scores than celebrities or individuals with great reputation. These results imply that the commercial measures should review their algorithms to avoid being influenced by automatic activity. Finally, del Campo-Ávila [2013] compared the scores of *Klout*, *PeerIndex* and *TwitterGrader*. They found that the *TwitterGrader* is not highly correlated with the other two metrics, whereas *Klout* and *PeerIndex* are highly correlated, as a result the features considered for measuring the influence of users must vary for each metric. Furthermore, the authors stated that, unlike *Klout* and *PeerIndex*, *TwitterGrader* is mainly focused on network topology.

6 Conclusions

This work aimed at shedding some light on how to assess user influence by proposing a novel metric based on user behaviour regarding both content-based and topological factors. The metric does not only consider each user individually, but also aims to assess the interactions with their neighbourhood.

The novel metric was compared to traditional means for assessing user influence, commercial metrics and a human assessment of user influence. The performed data analysis showed that there is no consensus on how the scores are computed and that defining user influence cannot be considered a trivial task. For example, in most cases, the commercial metrics proved to be independent from each other. Furthermore, results seemed to indicate that even among the different topics there might not be an uniform consensus regarding what means to be influential, which further remarks the fact that there is no unique definition of user influence, and that such definition might differ according to the analysed topic. Interestingly, the presented metric achieved its highest correlations with the human assessment of influence, which might indicate that the factors considered by the IS are truly relevant for assessing people's influence. Finally, results confirmed that only analysing the topological factors is not sufficient for accurately assessing the influence of users. Instead, an accurate assessment of user influence might also consider the published content and its influence over the neighbourhood of users.

Future work aims at analysing the influence of *Twitter* users taking into consideration the topics they post about. Furthermore, an extensive data analyses involving more *Twitter* users and human volunteers should be performed in order to obtain more statistical support for the reported results.

References

[Cha *et al.*, 2010] M. Cha, H. Haddadi, F. Benevenuto, and K.P. Gummadi. Measuring user influence in twitter: The million follower fallacy. In *4th International AAAI Conference on Weblogs and Social Media (ICWSM)*, 2010.

[del Campo-Ávila *et al.*, 2013] J. del Campo-Ávila, N. Moreno-Vergara, and M. Trella-López. Bridging the gap between the least and the most influential twitter users. *Procedia Computer Science*, 19(0):437 – 444, 2013. The 4th International Conference on Ambient Systems, Networks and Technologies (ANT 2013), the 3rd International Conference on Sustainable Energy Information Technology (SEIT-2013).

[Gaffney and Puschmann, 2012] Devin Gaffney and Cornelius Puschmann. Game or measurement? algorithmic transparency and the klout score. In *#influence12: Symposium & Workshop on Measuring Influence on Social Media Sep. 28-29*, volume 5, pages 1–2, 2012.

[Gladwell, 2000] Malcolm Gladwell. *The tipping point: how little things can make a big difference*. Little Brown, Boston, 1st edition, 2000.

[Java *et al.*, 2007] A. Java, X. Song, T. Finin, and B. Tseng. 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, San Jose, CA, USA, 2007.

[Katz and Lazarsfeld, 2005] Elihu Katz and Paul Lazarsfeld. *Personal Influence: The Part Played by People in the Flow of Mass Communications*. Transaction Publishers, October 2005.

[Keller and Berry, 2003] Edward Keller and Jonathan Berry. *The influentials: One American in ten tells the other nine how to vote, where to eat, and what to buy*. Simon and Schuster, 2003.

[Kendall, 1938] M. G. Kendall. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93, 1938.

[Kleinberg, 1999] Jon M Kleinberg. Authoritative sources in a hyperlinked environment. *Journal of the ACM (JACM)*, 46(5):604–632, 1999.

[Krishnamurthy *et al.*, 2008] B. Krishnamurthy, P. Gill, and M. Arlitt. A few chirps about Twitter. In *Proceedings of the 1st Workshop on Online Social Networks (WOSP'08)*, pages 19–24, Seattle, WA, USA, 2008.

[Kwak *et al.*, 2010] H. Kwak, C. Lee, H. Park, and S. Moon. What is Twitter, a social network or a news media? In *Proceedings of the 19th International Conference on World Wide Web (WWW'10)*, pages 591–600, Raleigh, NC, USA, 2010.

[Messias *et al.*, 2013] Johnnatan Messias, Lucas Schmidt, Ricardo Oliveira, and Fabrício Benevenuto. You followed my bot! transforming robots into influential users in twitter. *First Monday*, 18(7), 2013.

[Rogers, 2003] Everett M. Rogers. *Diffusion of innovations*. Free Press, New York, NY [u.a.], 5th edition, 08 2003.

[Tinati *et al.*, 2012] Ramine Tinati, Leslie Carr, Wendy Hall, and Jonny Bentwood. Identifying communicator roles in twitter. In Alain Mille, Fabien L. Gandon, Jacques Misselis, Michael Rabinovich, and Steffen Staab, editors, *WWW (Companion Volume)*, pages 1161–1168. ACM, 2012.

[Weng *et al.*, 2010] Jianshu Weng, Ee-Peng Lim, Jing Jiang, and Qi He. Twitterrank: Finding topic-sensitive influential twitterers. In *Proceedings of the 3rd ACM International Conference on Web Search and Data Mining (WSDM '10)*, pages 261–270, New York, NY, USA, 2010.

[Yamaguchi *et al.*, 2010] Yuto Yamaguchi, Tsubasa Takahashi, Toshiyuki Amagasa, and Hiroyuki Kitagawa. Turank: Twitter user ranking based on user-tweet graph analysis. In Lei Chen, Peter Triantafillou, and Torsten Suel, editors, *Web Information Systems Engineering WISE 2010*, volume 6488 of *Lecture Notes in Computer Science*, pages 240–253. Springer Berlin Heidelberg, 2010.