

De retibus socialibus et legibus momenti*

Daniel Gayo-Avello[†], David J. Brenes, Diego Fernández-Fernández,
María E. Fernández-Menéndez, Rodrigo García-Suárez

February 15, 2011

University of Oviedo (SPAIN) & Simplelogica, SL (SPAIN)

Abstract

Online Social Networks (OSNs) are a cutting edge topic. Almost everybody –users, marketers, brands, companies, and researchers– is approaching *OSNs* to better understand them and take advantage of their benefits. Maybe one of the key concepts underlying *OSNs* is that of influence which is highly related, although not entirely identical, to those of popularity and centrality. Influence is, according to Merriam-Webster, “*the capacity of causing an effect in indirect or intangible ways*”. Hence, in the context of *OSNs*, it has been proposed to analyze the clicks received by promoted *URLs* in order to check for any positive correlation between the number of visits and different “influence” scores. Such an evaluation methodology is used in this paper to compare a number of those techniques with a new method firstly described here. That new method is a simple and rather elegant solution which tackles with *influence* in *OSNs* by applying a physical metaphor.

Introduction

This paper describes an eminently empirical study for which a number of experiments were conducted. *Twitter* was chosen for that purpose because it is relatively easy to obtain data from it in comparison to other services such as *Facebook*. For those experiments the *Twitter* dataset from [6] was used. That study completely describes the data assets but, still, a brief description appears at the end of this introductory section.

*On social networks and the laws of influence

[†]Correspondence to: Daniel Gayo-Avello, Department of Computer Science (University of Oviedo) – Edificio de Ciencias, C/Calvo Sotelo s/n 33007 Oviedo (SPAIN), dani@uniovi.es

Basic concepts

Twitter is a *microblogging* service which allows users to publish text messages of up to 140 characters (*tweets*) which are shown to other users subscribed to the author feed (*followers*). Unlike other *OSNs*, relationships in *Twitter* are asymmetrical and, thus, it must be distinguished between people reading a given author messages (the aforementioned *followers*), and those persons that author reads (*friends* or *followees*).

The user graph and eigenvector centrality

Therefore, *Twitter* can be represented as a directed graph and, hence, it is amenable to analysis by means of eigenvector centrality algorithms. The aim of such algorithms is to compute the centrality of a node within a network (i.e. a graph) starting from rather simple assumptions: (1) the centrality of a node depends on the respective centrality values of the nodes linking to it; (2) the more nodes linking to a given one, or the more central the few nodes linking to it, the more central that node will be; (3) centrality values for all of the nodes within the network are iteratively computed until the algorithm converges.

Nonetheless to say, a number of algorithms exist to compute one or another “flavor” of these centrality scores. The “power iteration” method to compute the eigenvalues and eigenvectors of a matrix M is one of such methods. *PageRank* [16], *HITS* [10], *TwitterRank* [19], or *TunkRank* [17] are other approaches to compute slightly different scores better adapted to the graph properties of the WWW or the *Twitter* user graph.

The interested reader is recommended to consult [6] for a deep study and comparison of such algorithms regarding their application to the *Twitter* user graph. Suffice to say here that centrality algorithms are sensitive to a common form of abuse in *Twitter* –the *follow-to-be-followed* pattern– and, thus, robust methods to compute centrality are needed, in addition to verifying whether centrality is actually related or not to the elusive *influence*.

A naïve approach to popularity and influence

The number of *followers* has been largely considered as equivalent to *popularity*. After all, it seems rather obvious that the more *followers* a user has got the more popular s/he is and, in fact, celebrities such as Lady Gaga or Britney Spears have got millions of *followers*.

Given this simple approach to *popularity*, many *Twitter* users have exploited a simple rule of etiquette to get massive audiences. In *Twitter*, it is considered good manners to *follow back* a new *follower* and, hence, some abusive users (such as spammers and aggressive marketers) tend to follow thousands of people to get *followers* in return.

Because of this behavior, the *followers/followees* ratio has also been used as a proxy measure for a user’s *influence*: those users with a ratio greater than 1

are “influential” while those with a ratio lesser than 1 are “uninfluential”; besides, the larger the ratio the more “influential” the user.

Users’ tweeting behavior

In addition to *following* behaviors, *Twitter* users also get involved in *tweeting* behaviors: thus, a *tweet* can be original content produced by the author or it can be non-original content; that is, the user is repeating a *tweet* by another user (*retweeting*, in *Twitter* parlance).

Because *retweeting* is a form of citation, some syntax to provide attribution is needed. To do that, the name of the mentioned user is prepended with an *at sign*.

Let’s suppose, for instance, that Alice had *tweeted* the message “Hello world!”. If Bob wanted to *retweet* it he just should have to post¹: “RT @alice: Hello world!”, where *RT* stands for *retweet*.

Of course, this *mention* syntax is not limited to *retweets* and, indeed, it can be used to address other users and get involved into conversations rather similar to those within *IRC* (*Internet Relay Chat*).

So, in short, users can *tweet*, *retweet*, *mention*, or combine all of them – e.g. *retweeting* some content while addressing it to a third party: “RT @alice: Hello world! (cc @carol)”. For a deeper understanding of the *tweeting* and *retweeting* behavior of users we highly recommend the work by boyd *et al.* [2].

Last, but not least, every *tweet* is timestamped and, therefore, it is possible to compute for every user his *tweeting* frequency, bursts of activity, idle periods, etc.

Other approaches –different from eigenvector centrality– to compute influence

Feature rich approaches

Thus, *Twitter* provides plentiful of user features: *i*) number of *followers* and *followees*; *ii*) number of *tweets*, *retweets*, and *mentions* –both produced and received; *iii*) total number of *tweets*, and average number of *tweets* per hour, day, or week, etc.

All of these features are being used in almost any conceivable combination to produce formulae to score *Twitter* users. Companies such as *Klout*, *PeerIndex*, *tweetreach*, or *Twitalyzer*² use them to compute *ad hoc* scores which, arguably, provide a glimpse into the *influence* or *authority* of a given *Twitter* user.

Needless to say, actual details for such scores are undisclosed. Nevertheless, the interested reader can consult the recent work by Pal and Counts [14] describing a method to (1) cluster *Twitter* users according to several features

¹It must be noted that users in *Twitter* do not necessarily use their real name as user name. For instance, Ashton Kutcher is `aplusk` and CNN Breaking News is `cnnbrk`.

²<http://klout.com>, <http://www.peerindex.net>, <http://tweetreach.com>, and <http://twitalyzer.com>.

–including those aforementioned, and (2) rank users within the found clusters to find topical experts.

Influence maximization and diffusion cascades

A different angle of approach has been inspired by the highly influential work by Domingos and Richardson [4], and Kempe *et al.* [?]. Simply put, these researchers have studied the way in which *influence* (e.g. related to purchase intention) *virally* spreads through users within a network, so a minimum set of *influential* users can be found (i.e. those that should be addressed by a marketer in order to maximize sales with minimum effort).

The so-called *diffusion cascade models* –which are highly related to this area of *influence maximization*– have been rather successfully applied to *Twitter* (e.g. [18, 5, 12]).

It should be noted, however, that finding the optimum for influence maximization is NP-hard and, thus, efficient approximate algorithms are used. In this sense, Java *et al.* [8] showed that *PageRank* can be a feasible and inexpensive solution. Therefore, in spite of being a different approach, eigenvector centrality seems to be a good approximation to influence maximization in *OSNs* for all practical purposes.

The Influence-Passivity method

Finally, two recent works by Huberman *et al.* [7], and Romero *et al.* [15] must be cited. The former revealed important differences between the *Twitter* “declared” user graph and the actual interaction graph which is, in some sense, “hidden”. The “declared” user graph is built from the *follower-followee* relationships between users. The interaction graph, instead, does not take into consideration all of these relationships but only those which also involve actual interactions (i.e. users *mentioning* each other). Such a graph is “hidden” because interactions are not part of the user graph but have, instead, to be inferred from the *tweet* streamline.

The implications of this are clear: first, the number of *followers* and *followees* are misleading if there are no actual interaction between users; second, centrality measures obtained from the “declared” user graph could be very different from those obtained from the “hidden” interaction graph.

The second work is highly related to the first one; in it Romero *et al.* described the so-called *Influence-Passivity* algorithm. In certain sense this new algorithm is closely related to others such as *PageRank*, *HITS*, or *TunkRank*. However, unlike them, the edges (their weights, indeed) and partial scores are inferred from user interactions, in concrete, *retweets*.

The assumptions underlying their approach are very appealing: (1) The *influence* of a user depends on the *passivity* of his *followers* and, conversely, the *passivity* depends on the influence of his *followees*. (2) For each pair of users, an *acceptance* and a *rejection* rate are computed for the *follower* user. The former is the amount of *influence* the *follower* accepts from his *followee* (i.e. the number

of received messages s/he *retweets*) while the later is the amount of *influence* the *follower* *rejects*. (3) This way, the passivity of a user is proportional to both his *rejection rate* and the influence of his *followees*³ while the *influence* of a user is proportional to both the *acceptance rate* and the *passivity* of his *followers*. Finally, (4) *influence* and *passivity* scores are computed with an iterative algorithm that converges in relatively short time.

Dataset acquisition and description

The dataset used in this study consists of a collection of 27.9 million *tweets* and a user graph comprising 1.8 million users. Both were obtained using a number of methods of the *Twitter API* (*Application Programming Interface*). The *tweets* were collected from January 26 to August 31, 2009. Due to some network blackouts 4 days are missing and, thus, the dataset has got, on average, 130,000 *tweets* per day. On 2009 *Twitter* received 2.5 million *tweets* per day [18], hence, the data corresponds to about 5.6% of the total amount of *tweets* published during the crawling period.

Tweets are associated with metadata such as the publishing author and, thus, a list of 4.98 million users was obtained from the previous dataset and used to crawl the user graph. At the moment of that second crawling many accounts had been *suspended* or changed their status from public to private. Additionally, users without links to other users in the list were considered isolated and removed, and there were also minor network blackouts. For these reasons the graph contains less users than those publishing *tweets*. Anyhow, it was checked that the crawling was uniform and, in fact, the graph corresponds to 4% of the *Twitter*'s worldwide user base of 44.5 million users as of mid-2009 [16].

A proposal for Twitter dynamics

Rationale

It is clear that to apply any of the above-mentioned methods to compute *influence* in *Twitter* the user graph is needed. That graph alone is enough for eigenvector centrality methods but for the rest of approaches the published *tweets* are also required. Such data is needed to find out the *retweets*, *mentions*, diffusion cascades, and "hidden" relations between users.

Thus, researchers and practitioners working with *Twitter* usually deal with both data assets. It should be noted, however, that these two kinds of data (*tweets* and user graph) are not only distinct in nature but they are also crawled in very different ways.

Tweets can be relatively easy obtained as a data stream and most of the computation on them can be performed in near real-time. The user graph, however,

³That is, a user rejecting *tweets* from more influential users is more *passive* than a user rejecting the same amount but from less influential *followees*.

is a *snapshot* taken at a given time or, at most, a series of periodic *snapshots*. Nevertheless, *Twitter* in particular, and *OSNs* in general are highly dynamic systems, with users joining and quitting the network, and linking and unlinking among them continuously. Thus, static *snapshots* are a pale approximation to the actual evolving network.

Of course, it can be argued that a reasonable approximation is better than no approximation at all; however, in the light of recent findings such as those by Huberman *et al.* [7], we should wonder: *Is the user graph really needed to get a picture of Twitter? Even more concretely, is there any way of inferring influence by just relying on the most basic actions of Twitter users?*

The method described in the following subsection demonstrates that the user graph can be greatly disregarded, and *mentions* are enough to provide not only an accurate picture of *Twitter* but a dynamic one. Given that *mentions* are citations this should be hardly unsurprising; however, our approach is not based on bibliometrics but on physics, concretely on dynamic friction and uniformly accelerated linear motion.

A physical metaphor for influence in Twitter

The approach here described is an answer to the two aforementioned research questions and it evolved after a number of iterations.

Firstly, the role of the user graph to determine user *influence* was debated: the user graph could be (a) an essential data asset as in the cases of *PageRank* and *TunkRank*; (b) a starting point to find out the actual interaction graph as in the work by Huberman *et al.*; or (c) a dispensable asset. It is rather obvious that in the later case user interactions would be the only data to work on –no matter whether or not there were any *follower-follower* relationship between users.

Still, it was possible to build an implicit user graph from such users interactions in such a way that any graph-based method could be applicable. However, not building such an implicit graph was not only a novel approach but, besides, it would make real-time computations easier. Therefore, it was decided to study such an approach.

By totally disregarding both explicit and implicit user graphs it was clear that user influence would mostly rely on the *mentions* received by the users. It was also clear that a mere accounting of the total number of *mentions* received could be as misleading as the *follower* count. That is, it would not take account of the dynamic nature of *mentions* as they are related to events in which the mentioned user is involved. Hence, the time factor should be accounted for and, that way, it was perceived that equally or more important than the number of *mentions* would be the rate with which *mentions* increase. It was in that moment that the similarities to the dynamics equations were noted, and it was decided to study the feasibility of adapting them to compute users' *influence* in *OSNs*, concretely in *Twitter*.

Hence, in short, to devise this new approach, concepts from dynamics such as

force, *mass*, *acceleration* and *velocity* have been translated to an *OSN* scenario⁴. Thus, a user’s *influence* is modeled as *velocity* and, thus, *acceleration* can be used to detect trending users in real time.

Let’s start with Newton’s second law:

$$F = m \cdot a$$

How does this translate to *Twitter*? First, the *mass* of a user is the number of *followers*. Second, the *force* applied to put a user “in motion” is the number of *mentions* received (*retweets* alone could also be used). This way, a user with a high number of *followers* needs more *mentions* to start “moving” while a “lighter” user (one with a lower number of *followers*) requires fewer *mentions*.

It should be noted, however, that this equation assumes instantaneous *forces* and *accelerations*, and continuous time. For implementation purposes it is much more simple, however, to operate in discrete time intervals. Therefore, all of the experiments described in this paper were performed using one-hour sampling intervals. This way, the *force* applied on a user is, indeed, the number of *mentions* addressing that user in a given hour⁵.

In addition to this, under real circumstances there are more forces at stake: mainly, the force of kinetic friction F_f . Thus, *mentions* are actually the applied force, F_a , while F is the resultant force of F_a and F_f . The equation for the force of kinetic friction is the following:

$$F_f = \mu \cdot N$$

Where N is the normal force and μ the coefficient of friction. That way the *acceleration* would be:

$$a = \frac{F_a - \mu \cdot N}{m} = \frac{F_a - \mu \cdot m \cdot g \cdot \cos(\Theta)}{m} = \frac{F_a}{m} - \mu \cdot g \cdot \cos(\Theta)$$

Because the equation is to be translated to a non-physical scenario it can be simplified by supposing that not only μ , and g , but also $\cos(\Theta)$ are constant for every run of the method; thus, *acceleration* in *Twitter* would actually be:

$$a = \frac{F_a}{m} - \zeta$$

⁴We’d like to say that this is the first time that such a physical approach is suggested for *OSNs*; however, on October 20, 2010 the so-called “*velocity and acceleration*” model was reported. It must be said that, unlike our method, such a model is not an adaptation of physical laws but, instead, *velocity* and *acceleration* are used to denote the first and second derivatives of the time series corresponding to the tweet volume for a given topic [10]. Needless to say, both derivatives provide interesting information about the shape of the curve and, thus, the behavior of the topic but they are not a proper physical model and, hence, both their method and ours are unrelated.

⁵That sampling interval was fixed after some proof of concept experiments. When using shorter intervals (e.g. one minute) most of the users did not receive any mention, and when receiving any the “applied force” was virtually negligible. Larger intervals (e.g. one day) solved that problem but they were too coarse-grained for events evolving along hours.

Where ζ is a damping constant which is responsible for the decay of users' *acceleration* and *velocity* when they do not receive any *mention*. Needless to say, the value for that constant must be empirically determined and should have the same dimensions as the quotient $\frac{F_a}{m}$, that is, *mentions* per hour per *follower*. Hence, ζ value would be the average number of *mentions* per hour per user, divided by the number of *followers* an average user has got in *Twitter*.

Finally, the *velocity* of a *Twitter* user would be computed according to the following equation:

$$v_t = v_{t-1} + \frac{F_a}{m} - \zeta$$

It must be taken into account that (1) time is discrete, using one-hour sampling; (2) m is the number⁶ of *followers* of the user; (3) F_a is the number of *mentions* addressing the user in the last hour; (4) ζ is a constant positive number; and (5) negative velocities are not allowed and, hence, they should be replaced by zero.

From this equation it is easy to see that a frictionless scenario ($\zeta = 0$) is a special case where *velocity* is the accumulated number of *mentions* received by a user divided by his number of *followers*. Besides, if all of the users had the same number of *followers* then *velocity* would be equivalent to citation count.

In addition to that, by knowing both *velocity* and *acceleration* for each user at every hour it would be possible not only to know users *influence* but, much more importantly, to find *trending* users –i.e. those with higher *accelerations*– in real-time. Anecdotal evidence on this point is provided in a later section.

Experimental evaluation

Influence≈Attention≈Clicks

So far, another model to compute a score which may or many not relate to *influence* has been proposed. Thus, a way to correlate *velocity* with *influence* was also needed.

As it has been said, *influence* should exert measurable effects and, in this sense, the evaluation approach by Romero *et al.* [15] is pretty clever: they argued that, in the context of *Twitter*, *influence* should correlate with attention and, therefore, *URLs* posted by influential *Twitter* users should receive more visits than those *URLs* published by less influential ones.

Needless to say, the number of visits a given *URL* receives is just known to each website administrator. However, because of the length limit of the *tweets* (140 characters at most) virtually every published *URL* is *shortened* by means of one of several services⁷.

⁶In fact, smoother results can be obtained by applying natural logarithms to the number of *followers*.

⁷Using a shortening service a *URL* such as http://en.wikipedia.org/wiki/URL_shortening translates into <http://bit.ly/ebfVuu>.

One of them, *bit.ly*, provides an *API* which allows anyone to check the number of *clicks* a given *short URL* has received. Hence, using that *API*, it was possible to associate to *bit.ly URLs* appearing in *tweets* from the aforementioned dataset the corresponding number of visits those *URLs* had received⁸. Then, it was quite straightforward to check for any correlation between the *influence* of the users publishing the *URLs* and the visits for those *URLs*.

It must be noted, however, that some changes were made to the data collection methodology by Romero *et al.* They worked with *URLs* without taking into account for how long such *URLs* appeared in the *Twitter* stream. This is quite pertinent because some *URLs* can consistently appear for weeks or months, achieving a high number of visits with little or no relation at all with the *influence* of the users posting them⁹.

Therefore, in addition to preparing a *URL* dataset in the fashion of Romero *et al.*, a second one comprised of *URLs* appearing during one single week was prepared. An additional advantage of this second dataset is that it made possible to correlate *URL* visits with *velocity* values computed each week instead of comparing visits with one single final score for each user.

Finally, outliers were eliminated from both datasets using the common interquartile range method ($k = 1.5$). To that end, *URL* visits were considered and those *URLs* with exceedingly high numbers of visits were removed. In the second dataset, the outliers were computed for each different week and not for the whole dataset.

Hence, the first dataset was finally comprised of 22,920 *URLs* while the second one contained 10,120 *URLs* distributed in 29 weeks –from January 26 to August 16, 2009– with an average of 349 *URLs* and a standard deviation of 139.4.

Influence metrics evaluated

Romero *et al.* compared the predictive power of their *Influence-Passivity (I-P)* score with both *PageRank* and the number of *followers*. For this study not only those metrics were compared but also the recently proposed *TunkRank*, and the new method described in this paper –i.e. *velocity*.

Therefore, the number of *followers*, *I-P*, *PageRank*, and *TunkRank* were determined for those users appearing in the *Twitter* dataset described in [6]. In addition to that, *velocities* were computed and those reached at the end of each week were stored.

Needless to say, it was not possible to compute all of the scores for every user in the dataset: (1) *PageRank* and *TunkRank* require graph data for the

⁸Not every *bit.ly URL* appearing in the dataset was used, only those which were published by at least 3 users for whom graph data (i.e. their *followers* and *followees*) was available and it was possible to compute the *I-P* score by means of the *Influence-Passivity* method.

⁹For instance, *URLs* such as <http://bit.ly/SXp2X> or <http://bit.ly/2MbrXo> appeared virtually every week in our dataset; as it can be easily checked they are horoscopes. It is obvious that these are not the only websites that can recurrently appear in the *Twitter* stream (think for instance of news, auctions, or music sites).

users. (2) *Velocity* requires the users to be mentioned in the *tweets*. And (3), *I-P* does not only require graph data but also that connected users are involved in *retweeting* behavior. Thus, only those users qualifying for all of the methods were considered for the experiments¹⁰.

That way it was possible to associate every *URL* with both a number of visits, and a list of users who had “promoted” that *URL* in *Twitter*. Those users, in turn, had known “influence” scores. Therefore, it was just needed to look for any significant correlation between the number of clicks and each of the scores. To that end, the scores for those users promoting each *URL* were accumulated and, thus, for each *URL* a number of clicks and a single total “influence” score were available.

Some caveats should be noted. Firstly, when correlating “influence” with clicks from the *URL* dataset which ignored week limitations, the velocities employed were those reached by users on August 16, 2009 no matter the date when the *URL* had been published. Certainly, this is rather unrealistic but consistent with the way in which the rest of scores were obtained: after all, *PageRank* and *TunkRank* were computed from a graph crawled well after August 16, and the *retweets* required to applied the *Influence-Passivity* method were found across the whole dataset instead of using just the *tweets* predating the *URLs*.

Secondly, a single empirically found damping factor ($0 < \zeta \ll 1$) was applied to compute *velocity*. Proof of concept experiments showed that dynamic damping (i.e. a constant computed for each week or day based on the *tweeting* behavior of users during that period) did not provide better correlation. The same experiments revealed that a frictionless scenario ($\zeta = 0$) also shows a positive correlation between influence and clicks; however, the correlation was much weaker than when using a positive damping factor and, thus, such a frictionless model was disregarded.

Experimental findings

Pearson’s r was employed to compare *URL* clicks with accumulated “influence”. Certainly, assuming a linear regression model between a given “influence” score and *URL* visits can be an oversimplification but, hopefully, it could shed some light on the relation between such scores and observable events and, besides, it would make the results of this study comparable to those obtained by Romero *et al.* who reported R^2 values.

Table 1 shows the results obtained when comparing the aforementioned “influence” scores with the visits received by *URLs* in the dataset ignoring weekly limitations. Coefficients are not too high but, still, they are significant because of the sample size (22,920 *URLs*). From those results, it seemed that all of the “influence” scores exhibit a positive correlation with *URL* visits; however, some intriguing questions arise.

First of all, the results greatly departed from those reported by Romero *et al.* In fact, the correlation found in this study is much lower than the one reported

¹⁰These meant about 12,000 users; the strict requirements of the *Influence-Passivity* method –i.e. to just consider connected users who also *retweet* each other– drive to a very sparse graph.

by those researchers. In addition to that, the predictive performance of scores such as number of *followers*, *PageRank*, and *I-P* seems to be different than the one found by them. According to Romero *et al.* $followers < PageRank < I-P$ while Table 1 shows that $I-P < PageRank < followers$.

Needless to say, such differences could be attributed to many factors: from the datasets themselves to the way in which scores were computed. Romero *et al.* crawled just *tweets* containing *URLs* while the dataset employed in this study contained any kind of *tweet*. While they computed *PageRank* from the *retweeting* graph, for this study it was computed in a “traditional” way: i.e. from the *followers* graph. Besides, the way in which *URL* outliers were considered in both studies could also have distorted the results. Finally, while they compared average scores of the users promoting a *URL* with its clicks, accumulated scores were used for these experiments¹¹.

All of this would just mean that deeper analyses are needed; nevertheless, the attentive reader might have noted that a positive correlation between these “influence” scores and *URL* visits is not that surprising but, instead, expected. Indeed, the correlation between the number of *followers* and the clicks received provides a clue.

Certainly, algorithms such as *PageRank*, *TunkRank*, or *Influence-Passivity* are devised in such a way that users with few *followers* can still achieve rather high scores provided those few *followers* are “influential”. However, this is not the norm but the exception: most of the users with a high score also have a large number of *followers*. Hence, if users with a high *PageRank*, *TunkRank*, *I-P*, or *velocity* score have lots of *followers*, it is not that strange that the *URLs* they promote receive more visits than those promoted by users with lower “influence” scores. After all, they have much larger audiences and, thus, more visits are to be expected.

Table 2 reveals that a highly significant positive correlation exist between the number of *followers* and the different “influence” scores. In other words, the accumulated number of *followers* for the *URLs* must be considered a confounding variable and, thus, the data must be corrected for it¹².

To that end, both clicks and the different “influence” scores must be divided by the accumulated number of *followers*, in other words, the expected *audience* for each *URL*. This way, it would be checked if there exists any correlation between the probability for a member of a given audience to visit a *URL* and the portion of the *URL* promoters’ influence that member is responsible for.

Table 3 shows the results obtained after correcting the data for *audience*. The results for Influence-Passivity are inconclusive because there are no significant correlation. The rest of “influence” scores –namely *PageRank*, *TunkRank*, and *velocity*– show significant positive correlations. *Velocity* seems to be the best predictor, followed by *TunkRank* and, then, *PageRank*.

It should be remembered that all of these results were obtained from the

¹¹During the aforementioned proof of concept experiments it was found that average scores were worse predictors than accumulated scores.

¹²Although not directly related to the topic of this paper we cannot fail to urge the reader to consult the recent paper by West *et al.* [20].

“Influence” score	Pearson’s r	R^2	Significance
Number of <i>followers</i>	0.26637	0.07095	$p \ll 0.001$
<i>Influence (I-P)</i>	0.03627	0.00132	$p \ll 0.001$
<i>PageRank</i>	0.22381	0.05009	$p \ll 0.001$
<i>TunkRank</i>	0.17416	0.03033	$p \ll 0.001$
<i>velocity</i>	0.21981	0.04832	$p \ll 0.001$

Table 1: Correlation between different “influence” scores and clicks received by *URLs* in the dataset ignoring weekly limitations (data was not corrected for audience). Correlation coefficients are not very high but, given the size of the sample –22,920 *URLs*, all of them are significant ($p \ll 0.001$).

“Influence” score	Pearson’s r	R^2	Significance
<i>Influence (I-P)</i>	0.27994	0.07836	$p \ll 0.001$
<i>PageRank</i>	0.87284	0.76185	$p \ll 0.001$
<i>TunkRank</i>	0.75930	0.57653	$p \ll 0.001$
<i>velocity</i>	0.55151	0.30417	$p \ll 0.001$

Table 2: Correlation between the accumulated number of *followers* and the rest of “influence” scores using the information in the dataset ignoring weekly limitations. As it can be seen there exists a significant ($p \ll 0.001$) positive correlation between the number of *followers* and the “influence” scores. *Influence-Passivity* seems to be the method less sensitive to the number of *followers* and *PageRank* the most sensitive.

first *URL* dataset which did not take into consideration weekly limitations and, because of that, *velocity* scores were those reached by users on August 16, 2009. Another set of results was obtained by using the second dataset, comprising *URLs* which appeared in one single week.

For those experiments, three different *velocity* scores were employed: (1) *velocities* reached on August 16, 2009; (2) *velocities* computed at the end of each week; and (3) *velocities* computed at the end of the prior week. It is easy to see that the third “flavor” is the closest one to a real-time application.

Needless to say, the correlation coefficients reported in Table 4 were obtained by averaging the coefficients found for each week (cf. Cramer & Howitt [3], p.40) while the significance was computed according to the average sample size (349 *URLs* per week). These results are pretty consistent with those of Table 3: the correlation between *Influence-Passivity* and clicks is again non-significant; the rest of scores exhibit a significant positive correlation with *URL* visits; and, again, *velocity* is the best predictor.

On a side note, *velocities* computed on the week when *URLs* were published are slightly better predictors than *velocities* computed the week before. This would be of course expected if velocity in *Twitter* was a valid proxy measure for influence.

“Influence” score	Pearson’s r	R^2	Significance
<i>Influence (I-P)</i>	-0.01021	0.00010	Non-significant
<i>PageRank</i>	0.04399	0.00194	$p \ll 0.001$
<i>TunkRank</i>	0.13550	0.01836	$p \ll 0.001$
<i>velocity</i>	0.26532	0.07039	$p \ll 0.001$

Table 3: Correlation between different “influence” scores and clicks received by *URLs* in the dataset ignoring weekly limitations after correcting for the confounding variable *audience* (i.e. scores and clicks were divided by the accumulated number of *followers* of the users promoting the *URLs*). All of the scores, except for *I-P*, show significant positive correlations.

“Influence” score	Pearson’s r	R^2	Significance
<i>Influence (I-P)</i>	0.06806	0.00463	Non-significant
<i>PageRank</i>	0.25418	0.06461	$p \ll 0.001$
<i>TunkRank</i>	0.29921	0.08952	$p \ll 0.001$
<i>velocity</i> (August 16)	0.35464	0.12577	$p \ll 0.001$
<i>velocity</i> (on week)	0.37735	0.14240	$p \ll 0.001$
<i>velocity</i> (prior week)	0.37437	0.14015	$p \ll 0.001$

Table 4: Average correlation coefficients between different “influence” scores and clicks received by *URLs* in the dataset with weekly limitations. Both clicks and scores were corrected for the confounding variable *audience*. Reported coefficients were obtained by averaging the coefficients computed for each week.

Case study –Real-time detection of trending users by using acceleration

Perhaps one of the most direct applications of the new method described in this paper is to detect *trending* users; that is, those users reaching high *velocities* and who can be of interest for an audience that is still unaware of them.

The most straightforward way of finding such users would be computing the difference between the users’ current *velocities* and their previous ones to, then, order them by decreasing *acceleration*.

Nevertheless, by doing this there exists the risk of obtaining many users with high *accelerations* in absolute terms but rather low, even irrelevant, in relative terms (that would be the case of the most popular users, for instance).

To avoid this problem those users with a relative increase in *velocity* below a certain threshold (e.g. 10%) could be filtered out, and then the remaining users would be ordered by decreasing *acceleration*. This method was applied to the *Twitter* dataset to obtain a list of trending users for each week from January 26 to August 16, 2009.

A thorough evaluation of the quality of those results was out of the scope of this study; still, an informal analysis of the top ranked *trending* users was

conducted. To that end, the *tweets* mentioning the top-5 *trending* users for each week were obtained, and the most common phrases within them were obtained. Those phrases and the name of the user –generally a celebrity– were used to query a search engine. From the obtained results it was possible in virtually all of the cases to determine one or more actual events involving the user, and explaining the sudden increase in *velocity*.

Tables 5 to 9 show a summary of that informal evaluation; as it can be seen, the results obtained by applying the technique proposed in this paper seem highly promising.

Conclusions

This paper has described a new method to compute *Twitter influence* based on a physical metaphor which has got a number of advantages over commonly applied techniques.

First, it does not rely on the *Twitter* user graph which is costly to crawl, just provides static *snapshots* of a rapidly evolving network, and does not represent actual user interactions. Instead, the new method just requires the streamline of *tweets* to detect user *mentions*.

Second, it can be applied in near real-time and provides a natural way to detect *trending* or emerging users. Some anecdotal evidence on the quality of this approach has been provided.

A number of experiments were conducted to check whether the new *velocity* score actually correlates with *influence*. Results from those experiments have been reported, revealing that most of the commonly applied scores such as the number of *followers*, or *PageRank*, and recently proposed ones such as *TunkRank*, or *Influence-Passivity*, certainly exhibit a positive correlation with website visits.

However, it has also been shown that the number of *followers* is a confounding variable which must be accounted for. Therefore, it is not the total number of visits and the different “influence” scores which have to be correlated but, instead, the probability of a user visiting a promoted *URL* and the proportion of the promoter’s influence a single user is responsible for.

After correcting the data for the audience, it was revealed that all of the “influence” scores except for one –namely, *PageRank*, *TunkRank*, and *velocity*– exhibit positive correlation with user clicks and, thus, with *influence* in the sense of “attention gathering”. *Velocity*, the score inferred by the method proposed in this paper, was by a large margin the best predictor of user clicks.

The only score not showing significant correlation was *Influence-Passivity*. There exist, however, a number of reasons for this inconclusive result. The main one is, in all probability, the sparseness of the *retweet* graph obtained from the dataset because of the strict requirements of the *Influence-Passivity* method (i.e. a user has not only to *follow* another one but *retweet* some of his messages).

Hence, this study makes a number of contributions. (1) It adds to the general understanding of the concept of *influence* in *OSNs* and its relation to

“attention gathering”; (2) it has exposed the caveat due to the confounding nature of *audience* in this scenario; (3) it has shown how centrality measures can be used as rather good predictors of *influence*; and (4) it has described a new method that outperforms them with regards to *influence* scoring, and that can be applied in real-time to rank users and to detect emerging “influentials”. In this sense, an interesting future line of work would be studying the feasibility of adapting this new model to *tweets* themselves to detect *trending topics* and compare its performance with *Twitter*’s own implementation.

References

- [1] E. Bakshy *et al.* “Identifying ‘Influencers’ on Twitter,” *Proceedings of the fourth ACM International Conference on Web Search and Data Mining*, 2011.
- [2] d. boyd *et al.* “Tweet, Tweet, Retweet: Conversational Aspects of Retweeting on Twitter,” *Proceedings of the Hawaii International Conference on System Sciences*, 2010.
- [3] D. Cramer & D.L. Howitt, *The SAGE Dictionary of Statistics*, SAGE Publications Ltd., London, 2004.
- [4] P. Domingos & M. Richardson. “Mining the network value of customers,” *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, 2001.
- [5] W. Galuba *et al.* “Outtweeting the Twitterers - Predicting Information Cascades in Microblogs,” *Proceedings of the 3rd conference on Online social networks*, 2010.
- [6] D. Gayo-Avello. Nepotistic Relationships in Twitter and their Impact on Rank Prestige Algorithms. Arxiv preprint. arXiv:1004.0816, 2010.
- [7] B.A. Huberman *et al.* “Social networks that matter: Twitter under the microscope,” *First Monday*, vol. 14, no. 1–5, 2009.
- [8] A. Java *et al.* “Modeling the spread of influence on the blogosphere,” *WWW 2006 Workshop on Weblogging Ecosystem: Aggregation, Analysis and Dynamics*, 2006.
- [9] D. Kempe *et al.* “Maximizing the spread of influence through a social network,” *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, 2003.
- [10] S. Kerho. Predicting the Growth of Your Social Network, 2010, Available at: <http://threeminds.organic.com/2010/12/predicting-the-growth-of-your-social-network-2.html>

- [11] J.M. Kleinberg, “Authoritative sources in a hyperlinked environment”, *Proceedings of the ninth annual ACM-SIAM symposium on Discrete algorithms*, 1998, pp. 668–677.
- [12] C. Lee *et al.* “Finding Influentials from Temporal Order of Information Adoption in Twitter,” *Proceedings of 19th World-Wide Web (WWW) Conference* (Poster), 2010.
- [13] L. Page *et al.* The PageRank Citation Ranking: Bringing Order to the Web, 1998, Available at: <http://dbpubs.stanford.edu/pub/1999-66> (accessed November 26, 2010)
- [14] A. Pal & S. Counts. “Identifying Topical Authorities in Microblogs,” *Proceedings of the fourth ACM International Conference on Web Search and Data Mining*, 2011.
- [15] D.M. Romero *et al.* Influence and Passivity in Social Media. Arxiv preprint. arXiv:1008.1253v1, 2010.
- [16] E. Schonfeld. Twitter Reaches 44.5 Million People Worldwide In June, 2009. Available at: <http://techcrunch.com/2009/08/03/twitter-reaches-445-million-people-worldwide-in-june-comscore/> (accessed February 14, 2011)
- [17] D. Tunkelang. A Twitter Analog to PageRank, 2009. Available at: <http://thenoisychannel.com/2009/01/13/a-twitter-analog-to-pagerank/> (accessed November 26, 2010)
- [18] K. Weil. Measuring tweets, 2010. Available at: <http://blog.twitter.com/2010/02/measuring-tweets.html> (accessed February 14, 2011)
- [19] J. Weng *et al.* “TwitterRank: Finding Topic-sensitive Influential Twitterers,” *Proceedings of the third ACM international conference on Web Search and Data Mining*, 2010, pp. 261–270.
- [20] J. West *et al.* Big Macs and Eigenfactor Scores: Don’t Let Correlation Coefficients Fool You. Arxiv preprint. arXiv:0911.1807v2, 2010.

Week	Twitter user	Real name	Explanation and most frequent phrases
Feb. 1, 2009	stephenfry	Stephen Fry (English actor, writer, comedian, TV presenter and film director)	Stephen Fry was to appear on February 2, 2009 at an Apple Store in London to present his new audiobook. apple store
Feb. 8, 2009	wossy	Jonathan Ross (English TV and radio presenter)	(1) Ross, host for the 2009 edition of the Bafta Awards held on February 8, 2009, asked his followers for a word to insert during the ceremony; the chosen word was "salad". (2) On February 6, 2009 Tom Jones and Anna Friel, among others, visited Friday Night with Jonathan Ross. word salad, use word, good luck baftas, bafta word, twitter word, tom jones, anna friel
Feb. 15, 2009	lancearmstrong	Lance Armstrong (American professional road racing cyclist)	Lance Armstrong's time-trial bike was stolen on February 14, 2009 before the first stage of the Tour of California. stolen tt bike, time trial bike, bike stolen, tour california
Feb. 22, 2009 Mar. 1, 2009	the_real_shaq	Shaquille O'Neal (American professional basketball player)	(1) Mentions to the All-Star Game of the last weekend. (2) On February 20, 2009 O'Neal suggested to all of his followers to introduce themselves because they are connected in the, Shaquille wording, "Twittertonia". star game, twittertonia, public come say hi, twittertonia connect, congrats mvp On February 24, 2009 Shaquille suggested his followers to meet him in a mall to get two tickets. fashion sq mall, touches gets 2 tickets
Mar. 8, 2009	iamdiddy	Sean Combs (American record producer, rapper, actor, and fashion designer)	<i>Unknown.</i> bad boy, positive energy, first time, god bless

Table 5: Top *trending* users (ranking at the first position) found during February and early March 2009. The dates reported are those of the last day in the corresponding week. The third column identifies the individual or company and provides a short description. The last column provides an explanation for that user being *trending* on that week plus a number of the most frequent phrases found in the *tweets* mentioning the user during that week. As it can be seen, most of the times the *tweets* dealt with the actual events in which the user was involved.

Week	Twitter user	Real name	Explanation and most frequent phrases
Mar. 15, 2009	theellenshow	Ellen DeGeneres (American stand-up comedian, TV host and actress)	Ellen DeGeneres joined Twitter on March 10, on March 11 she was to appear at the Jay Leno show and she made a public appeal to get followers. In fact, most of the tweets mentioning @theellenshow were retweets of her original one: <i>"tweet & call everyone you know & tell them to follow me- I want to see how many I can get by the time I'm on Leno tonight."</i> leno tonight, tell follow, see many, many can get, call everyone
Mar. 22, 2009	lancearmstrong	Lance Armstrong (American professional road racing cyclist)	(1) On March 23, 2009 Lance Armstrong broke his collarbone in a crash during a race in Spain and had to face surgery. (2) On March 17, 2009 Lance Armstrong was required by French anti-doping agency to provide a hair sample. good luck, get well soon, best wishes, anti-doping, hope ok, recovery just, luck surgery, broken clavicle
Mar. 29, 2009 Apr. 5, 2009	macheist	MacHeist (website reselling Mac OS X shareware)	On March 24, 2009 the MacHeist 3 bundle was revealed in a live show. bundle reveal show, 3 bundle, buy bundle On March 25, 2009 the MacHeist 3 Bundle was on sale featuring 12 popular Mac applications normally valued at over \$900 for just \$39. macheist 3 bundle, mac apps, just \$39, mac apps worth \$900+, 12 top mac apps
Apr. 12, 2009	joeymcintyre	Joey McIntyre (American singer-songwriter and actor, part of the band New Kids on the Block)	The @joeymcintyre account was created on April 9, 2009 so, probably, that's the reason for it's sudden popularity. Most of the topics seem to be related to the "Full Service" summer tour in which NKOTB were involved. summer tour, happy easter, full service, easter bunny
Apr. 19, 2009	jordanknight	Jordan Knight (American singer-songwriter, part of the band New Kids on the Block)	<i>(Tentative)</i> Jordan Knight joined Twitter on April 14, 2009 and started to be addressed by fans with the rest of members of NKOTB. dannywood, jonathanrknight, donniewahlberg dannywood, joeymcintyre donniewahlberg
May 3, 2009	jonasbrothers	Jonas Brothers (American pop boy band)	On April 30, 2009 it was announced that Jonas Brothers would be participating in a series of live web chats starting on May 7. may 7th, live web chat may, question jonaslive

Table 6: Top *trending* users for March, April and early May. No data is provided for the week ending on April 26, 2009 because the dataset lacks several days on that week and, hence, all of the users lost *velocity*.

Week	Twitter user	Real name	Explanation and most frequent phrases
May 10, 2009	jordanknight	Jordan Knight (American singer-songwriter, part of the band New Kids on the Block)	<i>(Tentative)</i> Jordan tweeted “ <i>Tink! is the imaginary sound of my eyelids springing open when I wake up</i> ”. today show, joeymcintyre dannywood, donniewahlberg jonathanrknight, jonathanrknight joeymcintyre, tink sound
May 17, 2009	onlinesystem	Online System (¿online marketer?)	<i>(Tentative)</i> The user seems to be an aggressive marketer promoting systems to earn money throw affiliate marketing, virtually all of the tweets seem to be users reporting the increase in followers they got using the user’s method. followers using twitter, new followers added, 20 new followers added
May 24, 2009 May 31, 2009	jonasbrothers	Jonas Brothers (American pop boy band)	(1) “Paranoid” was the first single from their then new album; the video premiered on May 23, 2009. (2) Jonas Brothers play a little role in the movie “Night at the Museum: Battle of the Smithsonian,” sequel to the film “Night at the Museum,” which was released in theaters on May 22, 2009. music video, night museum 2, music video paranoid On May 28, 2009 another live web chat with the Jonas Brothers was held. web chat, web chat may 28th, new album, new songs, night museum 2
Jun. 7, 2009	mileycyrus	Miley Cyrus (American actress and pop singer)	“The Climb,” performed by Miley Cyrus for “Hannah Montana: The Movie,” won at the 2009 MTV Movie Awards held on May 31, 2009 in the category “Best Song from a Movie”. hannah montana, mtv movie awards, best song, congrats award, song climb, best song movie, congratulations
Jun. 14, 2009	peterfacinelli	Peter Facinelli (American actor)	Peter Facinelli made a bet with Rob DeFranco that he could get 500,000 followers on Twitter by June 19. If Facinelli wasn’t able to win the bet he should have to give DeFranco his Twilight chair. However, if he won the bet DeFranco should have to walk down Hollywood Blvd. in a bikini singing “All the Single Ladies”. win bet, single ladies, 500,000 followers, rob defranco, next week, bikini dance

Table 7: Top *trending* users for May and mid June 2009.

Week	Twitter user	Real name	Explanation and most frequent phrases
Jun. 21, 2009	perezhilton	Perez Hilton (American blogger and TV personality)	On June 17, 2009 Hilton used Twitter to claim assault by the Black Eyed Peas member will.i.am and his security guards. call police, black eyed peas, assaulted will, security wards
Jun. 28, 2009	songzyuuup	Trey Songz fan page (Trey Songz is an American recording artist, producer and actor)	Trey Songz attended and performed at the BET Awards ceremony held on June 28, 2009. bet awards, love trey, good bet awards, loved performance
Jul. 5, 2009	mileycyrus	Miley Cyrus (American actress and pop singer)	(1) Miley Cyrus starred in "Hanna Montana: The Movie" which as of July 2009 was still on theaters. (2) Cyrus started shooting the movie "The Last Song" on June 15, 2009. hannah montana, hanna montana movie, last song
Jul. 12, 2009	songzyuuup	Trey Songz fan page (Trey Songz is an American recording artist, producer and actor)	On June 2009 Trey Songz released a mixtape titled "Anticipation" through his blog before releasing this third album. trey songz, anticipation album, listening anticipation, mixtape anticipation
Jul. 19, 2009	jordanknight	Jordan Knight (American singer-songwriter, part of the band New Kids on the Block)	A concert by New Kids on the Block was live webcasted on July 17, 2009. webcast, jordan girl, love u, thank u, full service, good luck, luv u, love ya, miss u
Jul. 26, 2009	myfabolouslife	Fabulous (American recording artist)	<i>(Tentative)</i> Fan comments about the official remix of "Throw It in the Bag" featuring rapper Drake. The remix was released on August 18, 2009. throw bag, throw bag remix, ft drake, remix official

Table 8: Top *trending* users for June and July, 2009.

Week	Twitter user	Real name	Explanation and most frequent phrases
Aug. 2, 2009	paulaabdul	Paula Abdul (American pop singer, record producer, dancer, actress and TV personality)	During the 2000s Paula Abdul acted as judge on the TV contest “American Idol.” On July 17, 2009 her manager announced that she’d leave the show if producers didn’t step up a new deal. It wasn’t until August 4, 2009 that Paula definitely that she wouldn’t return to “Idol.” In the mean time many followers tweet their support for Abdul using the hashtag <i>#keeppaula</i> . #keeppaula, will continue
Aug. 9, 2009	adamlambert	Adam Lambert (American singer, songwriter, and actor)	(1) (<i>Tentative</i>) NOH8 Campaign was a silent protest photo project against California Proposition 8; it seems that Lambert fans were campaigning to get their idol taking part of the project. (2) On August 9, 2009 Adam Lambert won a Teen Choice Award. wewant 4noh8, noh8campaign, noh8campaign wewant, teen choice awards
Aug. 16, 2009			(1) On August 13, 2009 Lambert answered fan questions by means of Twitter in a so-called “Twitter party.” (2) On August 9, 2009 creative director of ELLE tweet about having Adam Lambert, among others, to the creative photo shooting for the next edition of the magazine. (3) On August 9, 2009 Adam Lambert won a Teen Choice Award. twitter party, elle shoot, details elle shoot, teen choice awards

Table 9: Top *trending* users for August, 2009.