“Who’s out there?” Identifying and Ranking Lurkers in Social Networks

Andrea Tagarelli
DIMES, University of Calabria
87036 Arcavacata di Rende (CS) - Italy
Email: tagarelli@dimes.unical.it

Roberto Interdonato
DIMES, University of Calabria
87036 Arcavacata di Rende (CS) - Italy
Email: rinterdonato@dimes.unical.it

***Abstract*—The massive presence of silent members in online communities, the so-called *lurkers*, has long attracted the attention of researchers in social science, cognitive psychology, and computer-human interaction. However, the study of lurking phenomena represents an unexplored opportunity of research in data mining, information retrieval and related fields. In this paper, we take a first step towards the formal specification and analysis of lurking in social networks. Particularly, focusing on the network topology, we address the new problem of *lurker ranking* and propose the first centrality methods specifically conceived for ranking lurkers in social networks. Using Twitter and FriendFeed as cases in point, our methods’ performance was evaluated against data-driven rankings as well as existing centrality methods, including the classic PageRank and alpha-centrality. Empirical evidence has shown the significance of our lurker ranking approach, which substantially differs from other methods in effectively identifying and ranking lurkers.**

I. Introduction

The majority of members of online communities play a passive or silent role as individuals that do not readily contribute to the shared online space. Such individuals are often called *lurkers*, since they belong to a community but remain quite unnoticed while watching, reading or, in general, benefiting from others’ information or services without significantly giving back to the community.

Lurking characterization in online communities has been a controversial issue from a social science and computer-human interaction perspective. Since the early works on social motivations and implications of lurking [1], [2], one common perception of lurking is that based on the infrequency of active participation (e.g., posting) to the community life, but other definitions have been given under the hypotheses of free-riding, legitimate peripheral participation [3], individual information strategy of microlearning [4], and knowledge sharing barriers (e.g., interpersonal or technological barriers) [5]. In the realm of social networks (SNs), negative views of the lurkers have been however supplanted with a neutral or even marginally positive view. A neutral perception of lurkers is related to the fact that their silent presence is seen as harmless and reflects a subjective reticence (rather than malicious motivations) to contribute to the community wisdom; half of times, a lurker simply feels that gathering information by browsing is enough without the need of being further involved in the community [2]. However, lurking can be expected or even encouraged because it allows users (especially newcomers) to learn or improve their understanding of the etiquette of an online community before they can decide to provide a valuable contribution over time. As a matter of fact, lurkers make up “the audience” of a community—it has been estimated that at any point in time approximately 90% of community members may be lurkers [1], [2]—and as such, it represents the crowd to attract. Tailoring online advertising strategies to the lurkers’ behavioral profile with the ultimate objective of *de-lurking* those users hence becomes an attractive opportunity.

Surprisingly, despite the fact that lurking has been recognized and surveyed in social sciences, we are not aware of any previous study that addresses the problem of ranking lurkers in a SN. Note that, beyond the frequent yet trivial case of users that exhibit a peripheral unstructured membership, hidden forms of lurking are massively present in SNs, which make it challenging to mine lurkers in a SN. While lurking is hard to track from a personal dispositional viewpoint, it appears that ranking lurkers is still possible by handling the situational variables that are related to the network of relationships between members. Moreover, a well-founded principle of eigenvector centrality like that we adopt in this work, will enable the determination of each node’s lurking score in function of the lurking scores of the nodes that it is connected to, based on global graph properties of propagation and attenuation of the information flowing through the network.

One may notice that ranking influential people is clearly valuable as we naturally tend to follow leaders and learn from them, and conversely wonder *“why ranking lurkers?”*. We argue that scoring community members as lurkers, rather than limiting to solely recognize (potential or actual) lurkers, should be seen as essential to determine the contingencies in the network under which different lurking behaviors occur, and ultimately to aid devising both generic and ad-hoc de-lurking plans and strategies. In effect, ordering members by decreasing lurking score would enable to manage priority in de-lurking applications, to identify the sub-communities particularly affected by lurkers, and to define personalized triggers of active participation. For example, lurkers of a given sub-community developed around an entity of interest (e.g., a person, or theme) would welcome messages that highlight the key topics (a service that is already delivered to its users by Twitter, for example), social events that describe how to approach a discussion in a forum or to start off your own project in a collaboration network, or introduce the role of forum moderators or team leaders. Moreover, in order to alleviate information overload, which is recognized as a major negative factor for participation, various mechanisms of filtering (e.g., recommending threads of discussion, providing visual maps of the categories of activities) could be applied with the ultimate goal of revealing the lurker's value (i.e., ideas, opinions, expertise) to the community.

Contributions and scope of this paper. In this paper we take a first step towards mining lurkers in SNs. We scrutinize the concept of lurking in SNs to determine the essential criteria that can be taken as the basis for mining lurkers. We lay out a basic topology-driven lurking definition, and propose various formulations of the lurker ranking problem that rely on the different aspects of our topology-driven lurking concept. By resorting to classic link-analysis ranking algorithms, PageRank and alpha-centrality, we provide a complete specification of lurker ranking methods.

We conducted experiments on Twitter and FriendFeed networks, whereby the evaluation goal was twofold: (i) to assess whether, upon modeling the directed graph underlying a SN to make it aware of the information the nodes receive rather than they produce, PageRank and related methods would be suited to solve such a new ranking problem in SNs; and (ii) to demonstrate the ability of our approach in capturing lurking cases that are intuitive yet non-trivial. Quantitative and qualitative results have demonstrated the effectiveness of our lurker ranking approach, highlighting superior performance against PageRank, alpha-centrality and the Fair-Bets model, which conversely might fail to correctly identify and rank presumed lurkers.

We would like to point out that other aspects than the network topology could not be covered in this paper, whose goal is to provide first solutions to a new ranking problem in SN. Indeed, as for instance PageRank was originally conceived for ranking web pages regardless of their content but based solely on their location in the Web, analogously here we just needed to focus on handling the network's graph structure to rank lurkers. Nevertheless, at the end of this paper, we will raise a number of challenges that capture increasingly complex intuitions of lurking, thus drawing up future developments of the lurker ranking problem.

II. TOPOLOGY-DRIVEN LURKING

User interactions in a SN are typically modeled as influence relationships, whose varying strengths are used to determine and rank the influential users.[1] In effect, ranking methods, such as PageRank, follow the conventional model of *influence graph*, which implies that the more (or more relevant) incoming links a node has the more important it is; for example, translated to Twitter terms, the more followers a user has the more interesting his/her published tweets might be. Recall however that, while this appears a reasonable approach, accounting for the node's in-degree solely should not be considered as reliable when determining the node's importance score. Rather, a combination of information based on incoming as well as outgoing links really matters. In any case, PageRank and related methods cannot be directly applied to lurking analysis because they assume that links across users carry the meaning of node influence propagation, which is related to the *amount of information (number of walks) a node generates*. By contrast, lurking behaviors build on the *amount of information a node receives*; again, in Twitter terms, if user v follows user u, then v is benefiting from u's information (i.e., v is receiving u's tweets), whereby relationship is modeled as a link from u to v. Within this view, a key notion for modeling the mutual contribution from incoming and outgoing links appears to be the (weighted) *in/out-degree ratio*: in an influence-oriented graph, this would map to the follower-to-followee ratio of a user (so that, the higher this ratio, the higher the probability that the user is influential), whereas in a lurking-oriented graph, the strength of a user's lurking status would be proportional to her/his followee-to-follower ratio. It should be noted that the significance of leveraging the in/out-degree ratio for ranking purposes in SN was already identified in [7], although again to score the authority of nodes. Upon the in/out-degree ratio intuition, we now provide a basic definition of lurking which aims to lay out the essential hypotheses of a lurking status based solely on the topology information available in a SN.

[1]Generic social links, like closeness relationships (i.e., friendships, acquaintances), should not be treated as valid indicators of real user interaction, as studied in, e.g., [6].

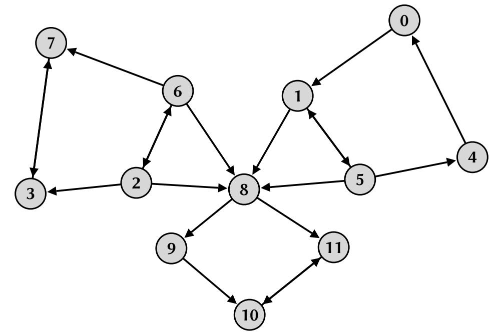

Fig. 1. An example SN graph

Definition 1 (Topology-driven lurking): Let $\mathcal{G} = \langle \mathcal{V}, \mathcal{E} \rangle$ denote the directed graph representing a SN, with set of nodes (members) $\mathcal{V}$ and set of edges $\mathcal{E}$, whereby the semantics of any edge (u, v) is that v is receiving information from u. A node v with infinite in/out-degree ratio (i.e., a sink node) is trivially regarded as a lurker. A node v with in/out-degree ratio above 1 shows a lurking status, whose strength is determined proportionally to (i) the in/out-degree ratio, (ii) the strength of non-lurking behavior shown by in-neighbors of v, and (iii) the strength of lurking behavior shown by out-neighbors of v.

The above definition states that determining a node's lurking behavior not only relies on its in/out-degree ratio but also on the extent to which its in-neighbors are rather influential nodes as well as its out-neighbors may in turn show a lurking behavior. To support this intuition, let us consider the example of network in Figure 1. Nodes $3, 7, 8, 10, 11$ have the highest in/out-degree ratio (i.e., 2), and as such they are candidate lurkers in the network. However, node 8 should be scored higher than others, since it benefits from information coming from two connected components, which are likely to contain influential nodes in the network (i.e., 5, 6). By contrast, nodes $10, 11$ should be scored as lurkers lower than node 8, since they are mainly fed by 8 itself; similarly, nodes $3, 7$ should be scored higher than $10, 11$ but lower than 8, since they receive information that propagates from a smaller subgraph.

III. LURKER RANKING

Given the directed graph $\mathcal{G} = \langle \mathcal{V}, \mathcal{E} \rangle$ representing a SN, for any node $v \in \mathcal{V}$ let $B(v) = \{u | (u, v) \in \mathcal{E}\}$ and $R(v) = \{u | (v, u) \in \mathcal{E}\}$ denote the set of in-neighbors (i.e., backward nodes) and out-neighbors (i.e., reference nodes) of v, respectively. The sizes of sets $B(v)$ and $R(v)$ are the in-degree and the out-degree of v, denoted as $in(v)$ and $out(v)$, respectively. Let $\mathbf{A}$ be the adjacency matrix of $\mathcal{G}$, with $A_{ij} = 1$ if $(v_i, v_j) \in \mathcal{E}$, and $A_{ij} = 0$ otherwise. Moreover, let $\widehat{\mathbf{D}_{in}}$ and $\widehat{\mathbf{D}_{out}}$ denote the in-degree and out-degree matrices, respectively. The in-degree matrix is a diagonal matrix defined as $\widehat{\mathbf{D}_{in}} = diag(\mathbf{e}^{\mathrm{T}}\mathbf{A})$, and the out-degree matrix is a diagonal matrix defined as $\widehat{\mathbf{D}_{out}} = diag(\mathbf{A}\mathbf{e})$, where $\mathbf{e}$ denotes a $|\mathcal{V}|$-dimensional column vector of ones. To deal with sink nodes and avoid infinite in/out-degree ratios, we introduce a Laplace smoothing factor in the definition of the diagonal matrices, and hence we will actually use the matrices defined as $\mathbf{D}_{in} = \widehat{\mathbf{D}_{in}} + \mathbf{I}$ and $\mathbf{D}_{out} = \widehat{\mathbf{D}_{out}} + \mathbf{I}$; consequently, $in(i)$ (resp. $out(i)$) is meant hereinafter as the actual in-degree (resp. out-degree) of node i plus one.

According to Definition 1, the simplest non-trivial form of a measure of lurker ranking is given by $r_i = \frac{in(i)}{out(i)}$, with $i \in \mathcal{V}$. However, the above form has clearly the disadvantage of assigning many nodes the same or very close ranks and, as we previously discussed, it ignores that the status of both the in-neighbors and out-neighbors contributes to the status of any given node. In the following we elaborate on each of those aspects separately.

In-neighbors-driven lurking: According to the second point in Definition 1, an in-neighbors-driven lurking measure can be defined as: $r_i = \sum_{j \in B(i)} \frac{out(j)}{in(j)} r_j$. Hence, the score of node i increases with the number of its in-neighbors and with their likelihood of being non-lurkers, which is expressed by a relatively high out/in-degree. The above formula however can be enhanced by including a factor that is inversely proportional to the i's out-degree:

$$r_i = \frac{1}{out(i)} \sum_{j \in B(i)} \frac{out(j)}{in(j)} r_j \tag{1}$$

Note that (1) accounts for both the contribution of a node's in-neighbors and its own in/out-degree property. Also, the set of equations defined by (1) has a matrix representation as follows: $\mathbf{r} = (\mathbf{D}_{out}{}^{-1}\ \mathbf{A}^{\mathrm{T}}\ \mathbf{D}_{out}\ \mathbf{D}_{in}{}^{-1})\ \mathbf{r}$.

Out-neighbors-driven lurking: The exclusive contribution of out-neighbors for the calculation of a node's lurking score, according to the third point of Definition 1, can be formalized as: $r_i = \sum_{j \in R(i)} \frac{in(j)}{out(j)} r_j$. However, this method would let the score of a node increase with the tendency of its out-neighbors of being lurkers, while ignoring the status of the node itself; as a consequence, not only reciprocal lurkers will be scored high but also every node from which lurkers receive information. A correction factor should hence be introduced as proportional to the in-degree of a node:

$$r_i = \frac{in(i)}{\sum_{j \in R(i)} in(j)} \sum_{j \in R(i)} \frac{in(j)}{out(j)} r_j \tag{2}$$

In (2), the in-degree of node i is divided by the sum of in-degrees of its out-neighbors in order to score i higher if it receives more than what its out-neighbors receive. The matrix form of (2) is as follows: $\mathbf{r} = (\mathbf{D}_{in}\ diag(\mathbf{A}\mathbf{D}_{in}\mathbf{e})^{-1}\ (\mathbf{A}\mathbf{D}_{in}\mathbf{D}_{out}{}^{-1}))\ \mathbf{r}$.

In-Out-neighbors-driven lurking: The previous definitions of in-neighbors-driven and out-neighbors-driven lurking can in principle be combined to obtain an integrated representation of all three aspects in Definition 1. Within this view, we define the score of node i as:

$$r_i = \left(\frac{1}{out(i)} \sum_{j \in B(i)} \frac{out(j)}{in(j)} r_j \right) \left(1 + \left(\frac{in(i)}{\sum_{j \in R(i)} in(j)} \sum_{j \in R(i)} \frac{in(j)}{out(j)} r_j \right) \right) \tag{3}$$

Note that in (3) we have emphasized the aspect related to the strength of non-lurking behavior of in-neighbors, which is expected to have a better fit of the hypothetical likelihood function for a given node.

A. PageRank and alpha-centrality based lurker ranking

The renowned Google's *PageRank* and the *alpha-centrality* methods will be used to provide a complete specification of the previously proposed models of lurker ranking. While being widely applied to a variety of application domains with the purpose of scoring the influence or prestige in information networks, the two methods rely on different assumptions which make it worth the exploration of lurker ranking through both approaches. In the following, we first recall the PageRank and alpha-centrality methods, then we will provide our PageRank- and alpha-centrality-based lurker ranking implementations.

PageRank and alpha-centrality in a nutshell: PageRank [8] extends the basic citation idea by assigning each page with a notion of importance that both relies on and influences the importance of neighboring pages. The PageRank vector is the unique solution of the iterative equation $\mathbf{r} = \alpha\mathbf{S}\mathbf{r} + (1-\alpha)\mathbf{v}$. $\mathbf{S}$ denotes the column-stochastic transition probability matrix, which is defined as $(\widehat{\mathbf{D}_{out}}^{-1}\mathbf{A})^{\mathrm{T}} + \mathbf{e}\mathbf{a}^{\mathrm{T}}/|\mathcal{V}|$, where $\mathbf{a}$ is a vector such that $a_i = 1$ if node i has zero out-degree, and 0 otherwise. Vector $\mathbf{v}$ is typically defined as $(1/|\mathcal{V}|)\mathbf{e}$, but can be modeled to bias the PageRank to boost a specific subset of nodes in the graph. Term α is a real-valued coefficient ($\alpha \in [0, 1]$, commonly set to 0.85), which acts as a damping factor so that the random surfer is expected to discontinue the chain with probability $1-\alpha$, and hence to randomly select a page each with relevance $1/|\mathcal{V}|$ (teleportation).

Alpha-centrality [9] expresses the centrality of a node as the number of paths linking it to other nodes, exponentially attenuated by their length. Moreover, it takes into account the possibility that each node's status may also depend on information that comes from outside the network or that may regard solely the member. Alpha-centrality is defined as $\mathbf{r} = \alpha\mathbf{A}^{\mathrm{T}}\mathbf{r} + \mathbf{v}$, where $\mathbf{v}$ is the vector of exogenous source of information ($\mathbf{v} = \mathbf{e}$ as default), and α here reflects the relative importance of endogenous versus exogenous factors in the determination of centrality. High values of α (e.g., 0.85) make the close neighborhood contribute less to the centrality of a given node. The rank obtained using alpha-centrality can

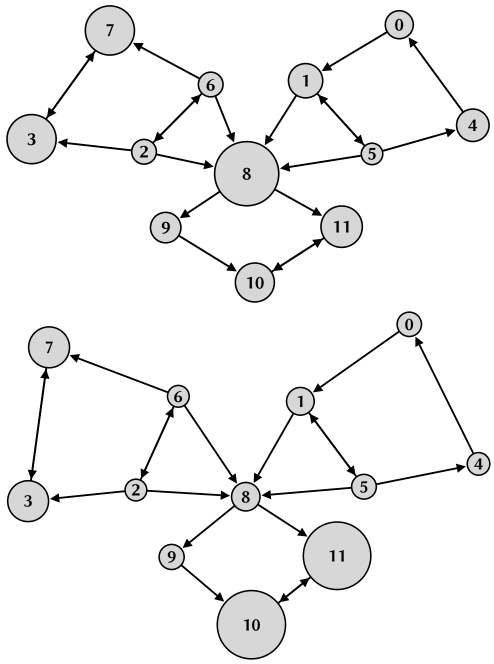

Fig. 2. Lurker ranking in the example SN graph of Fig. 1: LRin (on top) vs. PageRank (on bottom). Nodes are sized proportionally to their ranking scores.

be considered as the steady state distribution of an information spread process on a network, with probability α to transmit a message or influence along a link.

LurkerRank methods: We now define our lurker ranking methods upon the basic PageRank and alpha-centrality. We will refer to any of these algorithms as *LurkerRank* (for short LR), with prefix ac- to distinguish the alpha-centrality-based variants, and the various formulations of lurking with different suffixes in such a way that:

- LRin for the in-neighbors-driven LurkerRank:

$$r_i = \alpha \left(\frac{1}{out(i)} \sum_{j \in B(i)} w(j,i) \frac{out(j)}{in(j)} r_j \right) + \frac{1-\alpha}{|\mathcal{V}|} \tag{4}$$

- LRout for the out-neighbors-driven LurkerRank:

$$r_i = \alpha \left(\frac{in(i)}{\sum_{j \in R(i)} in(j)} \sum_{j \in R(i)} w(i,j) \frac{in(j)}{out(j)} r_j \right) + \frac{1-\alpha}{|\mathcal{V}|} \tag{5}$$

- LRin-out for the in-out-neighbors-driven LurkerRank:

$$r_i = \alpha \left[\left(\frac{1}{out(i)} \sum_{j \in B(i)} w(j,i) \frac{out(j)}{in(j)} r_j \right) \left(1 + \left(\frac{in(i)}{\sum_{j \in R(i)} in(j)} \sum_{j \in R(i)} w(i,j) \frac{in(j)}{out(j)} r_j \right) \right) \right] + \frac{1-\alpha}{|\mathcal{V}|} \tag{6}$$

TABLE I. EVALUATION NETWORK DATASETS

data	*# nodes*	*# links*	*avg in-degree*	*avg path length*	*# source nodes # sink nodes*	*avg in/out-degree* *
Twitter	16,009,364	132,290,000	8.26	5.91	1,067,936 10,298,788	2.65
FriendFeed	493,019	19,153,367	38.85	3.82	41,953 292,003	1.66

* Sink nodes and source nodes are excluded.

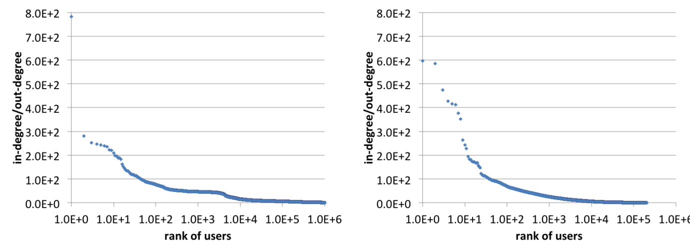

Fig. 3. In-degree to out-degree ratio distributions on *Twitter* (on the left) and *FriendFeed* (on the right). Source nodes and sink nodes are excluded.

- ac-LRin for the alpha-centrality-based in-neighbors-driven LurkerRank:

$$r_i = \alpha \left(\frac{1}{out(i)} \sum_{j \in B(i)} w(j,i) \frac{out(j)}{in(j)} r_j \right) + 1 \tag{7}$$

- and analogously ac-LRout and ac-LRin-out.

Note that in the LR formulations, we have introduced edge weights to deal with weighted graphs as well; although, as in our experimental setting, they are set as unitary by default. Figure 2 compares the rankings obtained by our LRin and basic PageRank on the example network of Figure 1 (α set to the default 0.85). Using LRin, node 8 was ranked highest (0.146), followed by 3 and 7 (0.112), and then 11 (0.094), 10 (0.088): this sheds light on the ability of LRin to match our definition of lurking (cf. discussion about Figure 1 in the previous section). By contrast, PageRank ranked first nodes 10 and 11 (both around 0.256), and then 3 and 7 with a significant gap in score from the first two (0.116), followed by 8 (0.052), 1 (0.048); moreover, node 5 was ranked eighth, despite it is a major feeder of the lurker 8, while it was correctly ranked lowest by LRin. Similarly, alpha-centrality (results not shown) did not fare well as it ranked first nodes 11 (0.317) and 10 (0.308), before ranking node 8 (0.095), and nodes 3 and 7 in ninth and tenth position both with a score of 0.004.

IV. Experimental Evaluation

A. Data

Two SNs were taken as cases in point for our evaluation, namely *Twitter* and *FriendFeed*. Beyond the complexity of their technical and sociological aspects, the two SNs have been selected since they naturally provide asymmetric relationships—recall that in our setting, a link from user i to user j means that j is a follower or subscriber of i—and also have quite different topological properties, as shown in Table I. Figure 3 also displays the in/out-degree distribution in each of the networks (due to massive presence of sink nodes which correspond to infinite in/out-degree, the figure concentrates on the nodes with finite in/out-degree).

From the Twitter dump studied in [10], we extracted the follower-followee topology starting from a connected component of one hundred thousands of users and their complete neighborhoods. This resulted in a network of about 16M users and 132M links. Actually, the building of our *Twitter* was constrained to the availability of a partial copy of the original dataset in [10] (we luckily got it before the Twitter's new Terms of Services was applied to avoid sharing tweet data) from which we gathered information about the number of retweets a user received. This information was exploited to define a Twitter-based ground-truth ranking and also to perform a qualitative evaluation on *Twitter*, as we shall describe later in this section.

FriendFeed is a real-time feed aggregator that aims to facilitate tracking users' social activities across multiple SN and bookmarking websites, blogs and microblogging services. We used the latest version of the dataset studied in [11], which offers information at various levels of interaction (subscription/like/comment). Note that the study pointed out the high similarity of the FriendFeed conversational/interaction scenario with Facebook (rather than with Twitter or other microblogging services). As suggested in [11], due to the recognized presence of spambots in this SN dataset, we filtered out users with an excessive number of posts (above 1200 posts in the monitoring period, or 20 posts per day).

B. Assessment methodology

Competing methods and notations: We compared our proposed methods against PageRank (henceforth PR), alpha-centrality (henceforth AC), and Fair-Bets model [7] (henceforth FB). The latter method, already mentioned in Section II, computes the score of any node i as $r_i = (1/out(i)) \sum_{j \in B(i)} r_j$. In addition, we included in the evaluation the *in/out-degree* distribution of the nodes in a network dataset, as a baseline method (henceforth IO).

Data-driven evaluation: Given the novelty of the problem at hand, we had to cope with an issue relating to the lack of ground-truth data for lurker ranking. In the attempt of simulating a ground-truth evaluation, we generated a *data-driven ranking* (henceforth DD) for each dataset, and used it to assess the proposed and competing methods.

On *Twitter*, we calculated the score of a node as directly proportional to its in/out-degree (Laplace add-one smoothed, cf. Section III) and inversely exponentially with a Twitter-specific measure of influence: $r_i^* = \frac{in(i)}{out(i)} \exp(-EI(i))$. $EI(\cdot)$ denotes the *empirical measure of influence* [12] which is used to estimate the influence of a twitterer based on the amount of information s/he posted (i.e., tweets) and that her/his followers have retweeted. For a twitterer i, $EI(i) = (1/out(i)) \sum_{j \in R(i)} nRetweets(j)$, where $nRetweets(j)$ is the number of retweets by follower j. Note that, as found in [10], a ranking based on retweets differs from that based on the number of followers, and this prompted us to combine the two aspects in our data-driven ranking.

We defined an analytically similar function for the *FriendFeed* data-driven ranking, in which the exponent (with negative sign) is the average add-one number of comments received by a user multiplied by the $\log_{10}$ of the add-ten number of posts by that user. Note that this combination of indicators of user's activity with user's influence was needed since only a limited portion (below 10%) of users in *FriendFeed* had information on the number of received comments.

Assessment criteria: In order to comparatively evaluate our proposed methods' performance w.r.t. the competing methods, we resorted to two well-known assessment criteria, namely *Fagin's intersection metric* [13] and *Bpref* [14].

Fagin measure allows for determining how well two ranking lists are in agreement with each other. This is regarded as the problem of comparing "partial rankings", since elements in one list may not be present in the other list. Moreover, according to [15], a ranking evaluation measure should consider top-weightedness, i.e., the top of the list gets higher weight than the tail. Applied to any two top-k lists $\mathcal{L}', \mathcal{L}''$, the Fagin score is defined as: $F(\mathcal{L}', \mathcal{L}'', k) = (1/k) \sum_{q=1..k} |\mathcal{L}'_{:q} \cap \mathcal{L}''_{:q}|/q$, where $\mathcal{L}_{:q}$ denotes the sets of nodes from the 1st to the qth position in the ranking. Therefore, F is the average over the sum of the weighted overlaps based on the first k nodes in both rankings.

Bpref [14] evaluates the performance from a different view, i.e., the number of non-relevant candidates. It computes a preference relation of whether judged relevant candidates R of a list $\mathcal{L}'$ are retrieved, i.e., occur in a list $\mathcal{L}''$, ahead of judged irrelevant candidates N, and is formulated as $Bpref(R, N) = (1/|R|) \sum_r (1 - (\#\text{of } n \text{ ranked higher than } r)/|R|)$, where r is a relevant retrieved candidate, and n is a member of the first $|R|$ irrelevant retrieved candidates. In our setting, we first determined N as the set of nodes with data-driven ranking score below or equal to 1, and used it for comparisons w.r.t. DD; whereas, for comparisons among competing methods, N was defined as the bottom of the corresponding method's ranking having the same size as N in the data-driven ranking. R was selected as the set of nodes having top-l% score from the complement of N.

Both F and $Bpref$ are within $[0, 1]$, whereby values closer to 1 correspond to better scores. For the experiments discussed in the following, we setup the size k of the top-ranked lists for Fagin evaluation to $k = 10^2, 10^3, 10^4$, and the l% of relevant candidates for Bpref evaluation to $l = 10, 25, 50$ (i.e., relevant candidates in the 90th percentile, the third quartile and the median). Moreover, unless otherwise specified, F scores will correspond to ranking lists without sink nodes, in order to avoid biasing (presumably overstating) our evaluation with trivial lurkers.

C. Results

1) Quantitative evaluation: We started our evaluation by first exploring the behavior of our LurkerRank methods (α set to the default value of 0.85) w.r.t. the data-driven ranking of each dataset (results not shown). On both datasets, we generally observed that, for any given setting of Fagin's k and Bpref's l parameters, the performance of LurkerRank methods, with the exception of LRout and ac-LRout, were not subject to significant fluctuations over the number of iterations. Moreover, it was interesting to observe on both datasets that all LurkerRank methods consistently reached a ranking stability very quickly, in the range 35÷75 iterations.

Tables II–III compare our LurkerRank methods against PageRank, alpha-centrality, Fair-Bets (all at convergence) as well as against DD and IO, for all variations of Fagin's and

TABLE II. COMPARATIVE PERFORMANCES ON *Twitter*.

	F $k = 10^2$ // 10^3 // 10^4					$Bpref$ $l = 10$ // 25 // 50				
	DD	IO	PR	AC	FB	DD	IO	PR	AC	FB
LRin	.527	.404	0.0	0.0	.112	.997	.992	.121	.790	.441
	.289	.209	0.0	0.0	.127	.995	.989	.473	.914	.704
	.581	.617	.001	.001	.068	.985	.962	.521	.866	.606
LRout	.030	.032	.181	.010	.034	.045	0.0	.754	.311	.313
	.008	.008	.351	.024	.015	.055	.001	.757	.650	.600
	.003	.002	.437	.048	.005	.109	.074	.641	.678	.648
LRin-out	.475	.364	0.0	0.0	.064	.968	.981	.039	.826	.204
	.314	.277	0.0	0.0	.063	.979	.977	.387	.929	.524
	.666	.688	.001	.001	.032	.961	.925	.453	.878	.489
ac-LRin	.583	.459	0.0	0.0	.174	.993	.990	.072	.808	.339
	.573	.570	0.0	0.0	.122	.992	.988	.443	.921	.653
	.767	.810	.001	.001	.048	.982	.967	.501	.872	.575
ac-LRout	.038	.032	.244	.006	.036	.049	0.0	.796	.339	.307
	.009	.008	.319	.017	.011	.059	0.0	.775	.659	.598
	.003	.002	.362	.042	.004	.120	.081	.654	.687	.643
ac-LRin-out	.473	.363	0.0	0.0	.062	.957	.981	.039	.828	.203
	.278	.234	0.0	0.0	.062	.975	.976	.386	.930	.464
	.663	.685	.001	.001	.031	.957	.933	.453	.880	.454

TABLE III. COMPARATIVE PERFORMANCES ON *FriendFeed*.

	F $k = 10^2$ // 10^3 // 10^4					$Bpref$ $l = 10$ // 25 // 50				
	DD	IO	PR	AC	FB	DD	IO	PR	AC	FB
LRin	.542	.690	.024	.010	.453	1.0	.980	.331	.606	.985
	.488	.586	.108	.118	.384	.998	.976	.570	.802	.977
	.576	.628	.126	.153	.493	.986	.953	.678	.843	.898
LRout	.015	.009	.479	.620	.011	.008	0.0	.691	.672	.031
	.138	.163	.550	.725	.167	.030	.038	.764	.746	.066
	.154	.156	.498	.704	.184	.062	.110	.739	.737	.258
LRin-out	.207	.297	.032	.042	.170	.972	.910	.252	.604	.879
	.278	.320	.061	.064	.166	.955	.910	.553	.794	.870
	.424	.455	.076	.099	.338	.914	.874	.642	.815	.813
ac-LRin	.575	.735	.025	.014	.467	1.0	.980	.300	.605	.980
	.520	.627	.118	.131	.403	.999	.977	.548	.803	.969
	.603	.660	.130	.161	.503	.988	.954	.661	.845	.882
ac-LRout	.015	.009	.479	.620	.011	.008	0.0	.691	.672	.031
	.138	.163	.550	.725	.167	.030	0.0	.749	.726	.066
	.154	.156	.498	.704	.184	.040	.080	.723	.718	.257
ac-LRin-out	.169	.243	0.0	0.0	.126	.958	.891	.237	.594	.852
	.240	.273	.001	.001	.122	.942	.892	.546	.785	.836
	.400	.426	.041	.064	.310	.898	.853	.634	.803	.782

Bpref's parameters. On *Twitter* (Table II), LRin and LRin-out along with their ac- counterparts showed a relatively much higher F intersection with DD (0.516 on average) and IO (0.473) than with FB (0.08), and a nearly empty F w.r.t. PR and AC. By contrast, LRout and ac-LRout exhibited a larger F with PR, although below 0.316 on average, while scoring even lower w.r.t. the other methods. Bpref evaluation led to mostly similar remarks on the relative comparison between proposed and other methods: LRin, LRin-out and their ac- counterparts highly matched DD and IO (around 0.97 on average), but also a moderately high $Bpref$ w.r.t. AC (0.87) and mid-low $Bpref$ w.r.t. FB (0.47). Again, like for the Fagin evaluation, LRout and ac-LRout showed no significant matches in practice with DD (while scoring pretty high w.r.t. PR): interestingly, this behavior confirms our intuition that determining the strength of lurking of a given node based on the strength of the lurking behavior shown by its out-neighbors (i.e., third aspect of Definition 1) is actually weaker than the other criteria given (and hence, other lurker ranking methods) and cannot provide meaningful results when used solely.

On *FriendFeed* (Table III), LRin and LRin-out along with their ac- counterparts were again the best-performing methods w.r.t. DD (0.42 F and 0.97 $Bpref$), and also showed mid F (0.34) and high $Bpref$ (0.89) w.r.t. FB. Yet, LRout and ac-LRout were moderately in agreement with PR and AC in terms of F, whereas all LR generally achieved mid $Bpref$ with both PR and AC.

We also determined the statistical significance of the better performance of LurkerRank methods w.r.t. the competing ones, through two stages of statistical testing analysis; in both cases,

TABLE IV. *Twitter* T-TEST ON THE PER-ITERATION PERFORMANCES.

	Fagin evaluation			*Bpref evaluation*		
	PR	AC	FB	PR	AC	FB
LRin	4.4E-65	4.4E-65	8.4E-11	5.2E-110	1.1E-25	2.1E-65
LRout	2.8E-41	2.7E-41	1.8E-04	3.2E-50	5.5E-79	9.2E-71
LRin-out	4.3E-277	4.4E-277	2.9E-12	1.5E-89	6.7E-21	7.6E-65
ac-LRin	5.6E-228	5.6E-228	4.8E-14	1.2E-91	2.1E-25	2.7E-65
ac-LRout	6.5E-34	6.2E-34	1.8E-04	4.1E-54	1.8E-71	2.3E-73
ac-LRin-out	3.8E-213	3.3E-265	3.4E-12	5.8E-85	2.1E-21	1.0E-64

TABLE V. *FriendFeed* T-TEST ON THE PER-ITERATION PERFORMANCES.

	Fagin evaluation			*Bpref evaluation*		
	PR	AC	FB	PR	AC	FB
LRin	1.3E-116	1.3E-103	2.6E-10	4.5E-195	5.9E-197	6.1E-10
LRout	8.5E-12	1.6E-101	1.5E-38	6.8E-252	1.3E-264	2.5E-271
LRin-out	6.0E-193	2.4E-166	2.1E-24	1.3E-298	2.1E-212	2.2E-116
ac-LRin	1.0E-195	1.0E-172	4.4E-13	5.0E-298	3.9E-189	7.8E-10
ac-LRout	2.6E-12	5.1E-88	1.3E-38	4.1E-99	5.9E-299	1.4E-282
ac-LRin-out	8.1E-63	1.3E-96	2.1E-25	8.3E-82	5.1E-226	1.5E-75

we fixed the Fagin parameter as $k = 10^4$ (which ensured a larger overlap between the ranking lists to be compared) and the Bpref parameter as $l = 25$ (for which $|R|$ was always smaller than $|N|$). Tables IV–V show the p-values resulting from an unpaired two-tail t-test, in which the performance scores obtained for each iteration by a ranking method w.r.t. DD were regarded as the statistical samples, under the null hypothesis of no difference in performance w.r.t. DD between a LurkerRank method and a competing method. Note that in all cases, the number of iterations (samples) was adequate to perform a t-test (generally above 50). Looking at the two tables and both F and $Bpref$ evaluation, the p-values turned out to be extremely low in most cases, thus giving a strong evidence that the null hypothesis was always rejected, at 1% significance level. This finding was useful to confirm that a certain difference (actually, the improvement) in performance between the LurkerRank methods and the competing ones, also on *FriendFeed* for which relatively high $Bpref$ scores were observed in the previous analysis.

In the second stage of statistical testing, we analogously performed a paired two-tail t-test in which the samples corresponded to the F scores respectively obtained by two ranking methods w.r.t. DD over the same randomly generated subgraph. For each of the network datasets, we extracted 100 subgraphs, each time starting from a randomly picked seed node and roughly covering a fixed number of nodes (around 1/100 of the original network size). This test was hence intended to stress the ranking methods performing over a pool of subnetworks having different characteristics from each other, and from the whole original network as well; for instance, on *Twitter*, the subnetworks had average path length mean of 2.52 (0.86 stdev), and in/out-degree ratio mean of 0.07 (0.13 stdev)—this might be explained because of the adopted approach of breadth-first traversal of the network, which led to connect the majority of nodes with a few source nodes having very high out-degree. On *Twitter*, we observed a close behavior between the LurkerRank methods (except LRout and ac-LRout) and AC (around 0.19 F on average), and between PR and FB, which however achieved a lower average F (0.029)—note that k was still set to 10^4, hence very high for such network sizes (i.e., around 200,000 nodes). In any case, i.e., for each pair of LurkerRank method vs. competing method, the null hypothesis of equal means was rejected even at 1% significance level, since the p-values were ranging from 1.4E-3 to 2.8E-19. Analogous final remarks were drawn for *FriendFeed*.

TABLE VI. TOP-20 *Twitter* USERS BY LURKING SCORE.

rank	PR			AC			FB			LRin		
	user	*score*	*#rt*	*user*	*score*	*#rt*	*user*	*score*	*#rt*	*user*	*score*	*#rt*
1	B.O.	4.85E-03	17811	B.O.	1.59E-04	17811	D.W.S.	1.15E-05	0	R.F.	7.78E-06	0
2	W.F.	3.57E-03	1676	ZAI.	1.41E-04	10902	n.a.	6.35E-06	0	R.J.	7.72E-06	0
3	ZAP.	2.47E-03	8707	ZAP.	1.36E-04	8707	APA.	6.33E-06	0	R.M.K.	7.49E-06	0
4	TH.	1.86E-03	7169	AS.	1.35E-04	1172	T.S.C.	5.41E-06	1	B.B.P.	7.35E-06	0
5	L.E.	5.77E-04	683	M.M.	1.31E-04	7	n.a.	5.07E-06	0	TR.	6.84E-06	0
6	J.B.	5.64E-04	1248	W.F.	1.23E-04	1676	CON.	4.97E-06	0	MU.	6.04E-06	0
7	M.S.	4.87E-04	476	M.K.	1.20E-04	48	K.T.	4.95E-06	0	B.R.	5.37E-06	0
8	AS.	4.25E-04	1172	P.B.	1.10E-04	328	n.a.	4.78E-06	0	AZ.	5.30E-06	0
9	OH.	3.53E-04	1009	W.A.	1.07E-04	2814	S.M.	4.36E-06	0	O.L.	5.25E-06	0
10	H.T.	3.19E-04	43	C.B.	1.06E-04	11943	n.a.	4.06E-06	0	N.T.	5.20E-06	0
11	E.T.	3.17E-04	2435	EL.	1.04E-04	902	n.a.	3.83E-06	0	FR.	5.15E-06	0
12	SCH.	3.02E-04	3277	SCO.	1.03E-04	6970	M.P.	3.82E-06	0	D.W.S.	5.15E-06	0
13	RE.	2.93E-04	1467	WI.	1.02E-04	811	n.a.	3.81E-06	0	AW.	4.96E-06	0
14	H.S.	2.89E-04	1346	O.W.	1.02E-04	1803	n.a.	3.79E-06	0	O.B.	4.68E-06	0
15	M.M.	2.89E-04	7	T.B.B.	9.84E-05	102	M.E.	3.69E-06	0	N.C.	4.56E-06	0
16	ZAI.	2.85E-04	10902	T.S.	9.82E-05	74	B.B.P.	3.68E-06	0	D.P.	4.43E-06	0
17	SCO.	2.84E-04	6970	S.S.	9.72E-05	789	n.a.	3.68E-06	0	AU.	4.30E-06	0
18	M.K.	2.63E-04	48	M.W.	9.17E-05	363	n.a.	3.61E-06	0	EM.	4.28E-06	0
19	WI.	2.59E-04	811	H.R.	8.89E-05	750	n.a.	3.58E-06	0	DI.	4.12E-06	0
20	W.A.	2.56E-04	2814	A.K.	8.69E-05	1572	n.a.	3.57E-06	0	M.A.	3.91E-06	0

For privacy reasons, users' names were replaced with their initials or abbreviations.

2) Qualitative evaluation: We investigated the meaningfulness of the rankings produced by LurkerRank methods, and compared them to those produced by PageRank, alpha-centrality, and Fair-Bets. Table VI reports the highest ranking 20 Twitter users for each method; additionally, the table also reports the number of times a user was retweeted (denoted as *#rt*). Among the LurkerRank methods, we selected LRin as it was consistently found as one of the best performing methods. Moreover, we left sink nodes out of consideration in order to avoid biasing our evaluation with trivial lurkers.

By comparing the top-ranked lists, it is evident that LRin behaved differently from the other algorithms, since it shared just two users with FB (dark-grey shaded) and no users at all with PR and AC. Interestingly, the LRin top-ranked list contains only users who have never been retweeted; furthermore, by retrieving the tweet post dates from Twitter, those users were all found as quite longer-time users, as in fact they joined Twitter much earlier (e.g., #8, #10 and #12 joined in 2007) than most users in the AC and PR top-ranked lists. Conversely, in the latter two lists most users have been significantly retweeted although they joined later (e.g., 2009).

PR and AC showed a certain association, with ten users in common (light-grey shaded). Most users in both the AC and the PR lists however were retweeted hundreds times, and hence they should not be considered as lurkers. Our hypothesis of non-lurking for those users was fully confirmed as we observed that those users' retweets were actually spread over a relatively short period of time (e.g., second half of 2009). Moreover, AC and PR ranked the same user on top, who is also the one having the highest number of retweets in the lists; indeed, that user is a very influential person, and in fact s/he has a followee/follower ratio much below 1: this would indicate that both AC and PR were not able to correctly handle this case (i.e., scoring it low enough), because their performance would be more affected by highly influential incoming links (i.e., followees)—which is a clear indication of tendency to absorb valuable knowledge—rather than by the number and type of followers. We also found other cases with characteristics similar to #1, e.g., #12 in the PR list, #10 and #14 in the AC list, and the common users "ZAP." (#3 in both lists) and "SCO." (#17 in PR, #12 in AC).

As concerns FB, it was surprising to find that 15 out of 20 top-ranked users actually refer to spammers (#4, a fashion/cosmetic marketing spammer, #9, in advertising, and #15, a porn spammer), or in general to suspended accounts (#2-3, #5, #8, #10-11, #13-14, #17-20). Only #6, #12 and #16 appear to be lurkers, which might be confirmed by their high in/out-degree ratio coupled with a zero retweet-count. By contrast, #1 is an art director and designer, and #7 refers to an account actively used for academic advising purposes; probably, the high number of followees (e.g., about 1800 for #7) has misled the method. Therefore, like PR and AC, FB might also fail to correctly recognize real lurkers.

Due to space limits, here we have presented only results on *Twitter*, but we ensure that similar findings were obtained on *FriendFeed* both in terms of evidence of the effectiveness of LRin and relatively poor reliability of the other methods.

V. RELATED WORK

The topic of lurking has been long studied in social science and recently has gained renewed interest in the computer-human interaction community. [16] investigates relations between lurking and cultural capital, i.e., a member's level of community-oriented knowledge. Cultural capital is found positively correlated with both the degree of active participation and, except for longer-time lurkers, with de-lurking. [17] leverages the significance of conceptualizing the lurking roles in relation to their boundary spanning and knowledge brokering activities across multiple community engagement spaces. The study proposed in [18] raises the opportunity of rethinking of the nature of lurking from a group learning perspective, whereby the engagement of intentional lurkers is considered within the collective knowledge construction activity. In [19], a set of statistical patterns is presented to characterize a comparison between contributing actors and lurkers across multiple communities, under the hypotheses of personal trait, engagement, and social learning. Exploring epistemological motivations behind lurking dynamics is the main focus of the study in [20], which indeed reviews major relevant literature on epistemic curiosity in the context of online communities and provides a set of propositions on the propensity to lurk and de-lurk. However, as with [17], the paper only offers insights that might be useful to guide an empirical evaluation of lurkers' emotional traits.

To the best of our knowledge, there has been no study other than ours that provides a formal computational methodology for lurker ranking. The study in [21], which aims to develop classification methods for the various SN actors (i.e., leaders, spammers, associates, and lurkers), actually treats the lurking problem marginally, and in fact lurking cases are left out of experimental evaluation. Similarly, [22] analyzes various factors that influence lifetime of SN users, also distinguishing between active and passive lifetime; however, analyzing passive lifetime is made possible only when the user's last login date is known, which is a rarely available information.

VI. CHALLENGES AND FUTURE WORK

The inherent complexity of lurking would advise that more information besides the network topology, certainly including temporal as well as contextualization aspects, might be considered for an improved ranking of lurkers. Further developments of the lurker ranking problem should hence take into account the aspects discussed below.

Time-driven lurking: Users in a SN naturally evolve over time playing different roles, thus showing a stronger or weaker tendency toward lurking on different times. Temporal dimension should be considered in terms of online access frequency of the community members. Lurkers tend to have unusual frequency of online presence, and hence any knowledge on the average central tendency/variability of online participation frequency of members in the community could guide the identification of critical time intervals to reveal lurking behaviors.

Context/Content-biased lurking: Lurking might also be identified w.r.t. a particular context, since the same user node can be involved in different aspects underlying the relationships in the network. Example contexts might refer to the type of resources exchanged, the topics discussed, the work tasks or interests shared among the community members. This would also suggest extensions of the lurking conceptualization to take into account the content semantics of the data associated with a selected context. For instance, by taking into account the content of the postings, including the messages sporadically posted by lurkers, the context-biased activity of a lurker would be better understood.

In their life cycle, users get different skills and experiences, and may change their interests. This clearly not also impacts on the role a user may take but also on the strength of the attachment the user has in relation to other users. Therefore, context and time dimensions might be jointly considered to determine which users may lurk. For instance, a lurking analysis of the context-biased activity of a user (e.g., number of postings) in function of her/his online access frequency would help to get a more complete picture of the user's engagement level in the community.

Boundary-spanning lurking: We believe that a further attractive option for conceptualizing lurking behaviors lies in the theory of boundary-spanning and knowledge transfer: given the very large scale of SNs and their heterogeneity in the type of information exchanged and actors, connected components are likely to be formed and dynamically change over time. Therefore, some of the members that lay on the boundary of a component may bridge over other components. As a result, members who lurk inside a component may not lurk, or even take on the role of experts, in other connected components.

It is also worth noting that combining topology with temporal, context and content information would also help understanding relations between lurkers and other SN actors. While naturally interpreted as an inverse notion of influence, lurking would actually seem quite unrelated to the other major event in a SN, i.e., spamming, and in effect, personal dispositions and behaviors are totally different for the two types of actors. However a lurker's life-cycle could be somehow affected by the presence of spammers in the SN. Hence we raise a question whether existing solutions to spam detection and trust analysis could support both the identification and ranking of lurkers and eventually their delurking.

VII. Conclusion

We addressed the previously unexplored problem of ranking lurkers in a SN. We introduced a topology-driven lurking definition and proposed various lurker ranking methods, for which we provided a complete specification in terms of the well-known PageRank and alpha-centrality. We have been positively impressed by results achieved on Twitter and FriendFeed by some of our lurker ranking methods, especially in terms of significance and higher meaningfulness w.r.t. other competing methods. Future directions of research have also been issued.

References

[1] B. Nonnecke and J. J. Preece, "Lurker demographics: counting the silent," in *Proc. ACM CHI*, 2000, pp. 73–80.

[2] J. J. Preece, B. Nonnecke, and D. Andrews, "The top five reasons for lurking: improving community experiences for everyone," *Computers in Human Behavior*, vol. 20, no. 2, pp. 201–223, 2004.

[3] J. Lave and E. Wenger, *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press, 1991.

[4] N. Kahnwald and T. Köhler, "Microlearning in virtual communities of practice? an explorative analysis of changing information behaviour," in *Proc. Microlearning*, 2006, pp. 157–172.

[5] A. Ardichvili, "Learning and knowledge sharing in virtual communities of practice: motivators, barriers, and enablers," *Advances in Developing Human Resources*, vol. 10, pp. 541–554, 2008.

[6] C. Wilson, A. Sala, K. P. N. Puttaswamy, and B. Y. Zhao, "Beyond Social Graphs: User Interactions in Online Social Networks and their Implications," *ACM Trans. on the Web*, vol. 6, no. 4, p. 17, 2012.

[7] S. Budalakoti and R. Bekkerman, "Bimodal invitation-navigation fair bets model for authority identification in a social network," in *Proc. ACM WWW*, 2012, pp. 709–718.

[8] S. Brin and L. Page, "The anatomy of a large-scale hypertextual Web search engine," *Computer Networks and ISDN Systems*, vol. 30, no. 1-7, pp. 107–117, 1998.

[9] P. Bonacich and P. Lloyd, "Eigenvector-like measures of centrality for asymmetric relations," *Social Networks*, vol. 23, pp. 191–201, 2001.

[10] H. Kwak, C. Lee, H. Park, and S. B. Moon, "What is Twitter, a social network or a news media?" in *Proc. ACM WWW*, 2010, pp. 591–600.

[11] F. Celli, F. M. L. D. Lascio, M. Magnani, B. Pacelli, and L. Rossi, "Social Network Data and Practices: The Case of Friendfeed," in *Proc. Social Computing, Behavioral Modeling, and Prediction*, 2010, pp. 346–353.

[12] E. Bakshy, J. M. Hofman, W. A. Mason, and D. J. Watts, "Everyone's an influencer: quantifying influence on Twitter," in *Proc. ACM WSDM*, 2011, pp. 65–74.

[13] R. Fagin, R. Kumar, and D. Sivakumar, "Comparing Top k Lists," *SIAM Journal on Discrete Mathematics*, vol. 17, no. 1, pp. 134–160, 2003.

[14] C. Buckley and E. M. Voorhees, "Retrieval evaluation with incomplete information," in *Proc. ACM SIGIR*, 2004, pp. 25–32.

[15] W. Webber, A. Moffat, and J. Zobel, "A similarity measure for indefinite rankings," *ACM Trans. Information Systems*, vol. 28, no. 4, p. 20, 2010.

[16] V. Soroka and S. Rafaeli, "Invisible participants: how cultural capital relates to lurking behavior," in *Proc. ACM WWW*, 2006, pp. 163–172.

[17] J. Cranefield, P. Yoong, and S. L. Huff, "Beyond Lurking: The Invisible Follower-Feeder In An Online Community Ecosystem," in *Proc. Pacific Asia Conf. on Information Systems (PACIS)*, 2011, p. 50.

[18] F. C. Chen and H.-M. Chang, "Do lurking learners contribute less?: a knowledge co-construction perspective," in *Proc. Communities and Technologies (C&T)*, 2011, pp. 169–178.

[19] M. Muller, "Lurking as personal trait or situational disposition: lurking and contributing in enterprise social media," in *Proc. ACM CSCW*, 2012, pp. 253–256.

[20] A. Schneider, G. von Krogh, and P. Jager, ""What's coming next?" Epistemic curiosity and lurking behavior in online communities," *Computers in Human Behavior*, vol. 29, pp. 293–303, 2013.

[21] M. Fazeen, R. Dantu, and P. Guturu, "Identification of leaders, lurkers, associates and spammers in a social network: context-dependent and context-independent approaches," *Social Network Analysis and Mining*, vol. 1, no. 3, pp. 241–254, 2011.

[22] J. Lang and S. F. Wu, "Social network user lifetime," *Social Network Analysis and Mining*, vol. April, 2012.