

On the Efficiency of the Information Networks in Social Media

Mahmoudreza Babaei
MPI-SWS, Germany

Przemyslaw Grabowicz
MPI-SWS, Germany

Isabel Valera
MPI-SWS, Germany

Krishna P. Gummadi
MPI-SWS, Germany

Manuel
Gomez-Rodriguez
MPI-SWS, Germany

ABSTRACT

Social media sites are information marketplaces, where users produce and consume a wide variety of information and ideas. In these sites, users typically choose their *information sources*, which in turn determine what specific information they receive, how much information they receive and how quickly this information is shown to them. In this context, a natural question that arises is how *efficient* are social media users at selecting their information sources.

In this work, we propose a computational framework to quantify user's efficiency at selecting information sources. Our framework is based on the assumption that the goal of users is to acquire a set of unique pieces of information. To quantify user's efficiency, we ask if the user could have acquired the same pieces of information from another set of sources *more efficiently*. We define three different notions of efficiency – link, in-flow, and delay – corresponding to the number of sources the user *follows*, the amount of (redundant) information she acquires and the delay with which she receives the information. Our definitions of efficiency are general and applicable to any social media system with an underlying information network, in which every user follows others to receive the information they produce.

In our experiments, we measure the efficiency of Twitter users at acquiring different types of information. We find that Twitter users exhibit sub-optimal efficiency across the three notions of efficiency, although they tend to be more efficient at acquiring non-popular pieces of information than they are at acquiring popular pieces of information. We then show that this lack of efficiency is a consequence of the triadic closure mechanism by which users typically discover and follow other users in social media. Thus, our study reveals a tradeoff between the efficiency and discoverability of information sources. Finally, we develop a heuristic algorithm that enables users to be significantly more efficient at acquiring the same unique pieces of information.

Categories and Subject Descriptors: H.1.2 [Information Systems]: Models and Principles – *Human information processing*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WSDM'16, February 22–25, 2016, San Francisco, CA, USA.

© 2016 ACM. ISBN 978-1-4503-3716-8/16/02...\$15.00

DOI: <http://dx.doi.org/10.1145/2835776.2835826>

General Terms: Human Factors; Measurement; Performance.

Keywords: efficiency, information network, social media, information, cover set, rewiring algorithm, optimization, lossless.

1. INTRODUCTION

Over the last decade, the advent of social media has profoundly changed the way people produce and consume information online. A characteristic feature of many social media sites (e.g., Twitter or Pinterest) that distinguishes them from mainstream news media sites (e.g., CNN.com or NYTimes.com) is the *information network* created by consumers *following* their preferred producers of information [6, 18, 10]. However, the task of selecting information sources from potentially tens to hundreds of millions of users poses serious challenges and raises important questions that have not yet been addressed. For example, recent studies have observed that out of fear of missing out on important information, users tend to follow too many other users [14]. In the process, they receive a lot of redundant information [4], become overloaded, and effectively miss the information they are interested in [12, 19]. Moreover, it is very hard to ascertain the quality, relevance, and credibility of information produced by social media users [2, 8, 11]. Also, many users rely on their network neighborhood for discovering new sources of information, as observed by the large number of triadic closures [26, 13, 24] in link creation.

The motivation behind this work originates from two fundamental questions: 1) How *efficient* are the users of a social media site at selecting which other users to follow to acquire information of their interest? and, 2) Can we propose methods to enable a user to acquire the same pieces of information from another set of users in the social media site *more efficiently*? To answer these intertwined questions, we view the structure of the information networks in social media sites as the outcome of a network formation game [17], where a node (*i.e.*, user) links to other nodes to solve a specific task (*i.e.*, acquire information relevant to the user). In this work, we propose a general computational framework to quantify and optimize the efficiency of links created by users to acquire information.

With such a framework in place, we analyze the efficiency of information networks of social media users, addressing several important additional questions. For instance, a user might aim to receive pieces of information of different popularity. However, can popular information be covered as efficiently as non-popular information? Also, studies have identified triadic closure [26, 13, 24] and information diffusion [28, 23, 3] as salient mechanisms that trigger new link creation in social networks. However, how efficient are information networks created with these mechanisms? Similarly, it has been observed that in information systems such as

Wikipedia [29, 30], links are primarily established to make related content more easily discoverable, rather than for some nuanced notion of efficiency. Is the link creation in user-generated information networks such as Twitter driven by a similar goal of discovering users in the network neighborhood, rather than for efficiency in acquiring information?

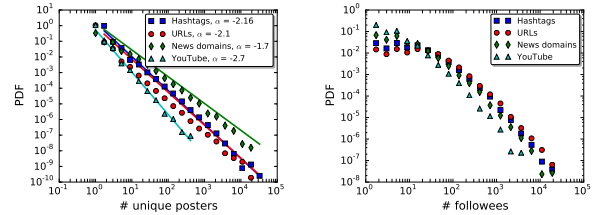
Our computational framework is based on the following key concept: given a set of unique ideas, pieces of information, or more generally, *memes* \mathcal{I} spreading through an information network, there is an *optimal* set of nodes that, if followed, would enable us to get to know \mathcal{I} . Naturally, this concept relies on defining what is an optimal set. Here, we consider three notions of optimality, which lead to three types of efficiencies:

- I. **Link efficiency.** The optimal set $\mathcal{U}^l(\mathcal{I})$ is the one that contains the smallest number of users. Then, we compute link efficiency by comparing the number of people a user follows, i.e., the number of *followers*, with the size of the optimal set $\mathcal{U}^l(\mathcal{I})$. Finding the optimal set reduces to a minimum set cover problem, which can be solved using a well-known and efficient greedy algorithm with provable guarantees [15].
- II. **In-flow efficiency.** The optimal set $\mathcal{U}^f(\mathcal{I})$ is the one that provides the least amount of tweets per time unit. Then, we compute in-flow efficiency by comparing the amount of tweets per time unit a user receives from the people she follows with the amount of tweets per time unit she would have received by following the users in the optimal set $\mathcal{U}^f(\mathcal{I})$. Finding the optimal set reduces to a minimum weighted set cover problem, which again can be solved efficiently with provable guarantees [15].
- III. **Delay efficiency.** The optimal set $\mathcal{U}^t(\mathcal{I})$ is the one that provides the memes as early as possible. Then, we compute delay efficiency by comparing the average delay per meme that the user achieves through the people she follows with the average delay she would achieve by following the users in the optimal set. Here, we define the delay at acquiring a meme in a social media system as the difference between the time when the user received the meme in her timeline and the time when the meme was first mentioned by a user in the social media system. Finding the optimal set reduces to finding the set of users who made the first mention of each of the memes in the social media system.

Although the concept of link efficiency has been previously introduced by us in [4], here we extend it to account for other definitions of efficiency, i.e., in-flow and delay efficiency. Another study related to ours uses set covers for efficient detection of outbreaks [21]. However, that related work does not compare the optimal sets to the sets encountered in reality.

In this study, we apply the three notions of efficiency to the information network of Twitter users. Our analysis not only shows how efficient users are at acquiring information, but also helps us in understanding the influence of different factors on efficiency:

1. We find that users acquire information sub-optimally with respect to the three notions of efficiency. However, the higher is the coverage of information that they want to acquire, the higher is the efficiency.
2. While the popular pieces of information are acquired inefficiently, less popular pieces of information are acquired more efficiently.



(A) Users per meme

(B) The number of followers

Figure 1: The distributions of (A) the number of users posting a unique meme and (B) the number of followers posting a specific type of meme at least once. For (A), a power-law is fitted and the exponent α is given.

3. Users trade followers’ discoverability for information efficiency. While, for a typical Twitter user, many people in her original ego-network may be discovered by triadic closure, very few people in the optimized efficient ego-networks can be discovered in that way.
4. We introduce a heuristic algorithm that considerably increases both user’s inflow and delay efficiencies and delivers to her the same unique pieces of information, by rewiring user’s information network.

Our empirical findings also shed light on how our notion of efficiency relates to user’s ego-network structure (*e.g.*, triangle closure, clustering coefficient).

2. DATASET

We use a large Twitter dataset, as reported in previous work [9], which comprises the following three types of information: profiles of 52 million users, 1.9 billion directed follow links among these users, and 1.7 billion public tweets posted by the collected users. The link information of the network is based on a snapshot taken at the time of data collection, in September 2009. In our work, we limit ourselves to tweets published during one week, from July 1, 2009 to July 7, 2009, and filter out users that did not tweet before July 1, in order to be able to consider the social graph to be *approximately* static. After this filtering, we have 395,093 active users, 39,382,666 directed edges, and 78,202,668 tweets.

Then, we sample 10,000 users at random out of the 395,093 active users and reconstruct their timelines by collecting all tweets published by the (active) people they follow (among all the 395,093 users), build their ego-networks (*i.e.*, who follows whom among the people they follow), and track all the unique memes they were exposed to during the observation period. We consider four different types of memes:

- I. **Hashtags.** Hashtags are words or phrases inside a tweet which are prefixed with the symbol “#”. They provide a way for a user to generate searchable metadata, keywords or tags, in order to describe her tweet, associate the tweet to a (trending) topic, or express an idea. Hashtags have become ubiquitous and are an integral aspect of the social Web nowadays [25].
- II. **URLs.** We extract all URLs mentioned inside tweets [22]. Since most of URLs in Twitter are shortened, we unwrap them by calling the API of the corresponding shortening service. Here, we considered seven popular URL shorteners:

bit.ly, tinyurl.com, is.gd, twurl.nl, snurl.com, doiop.com and eweri.com, and discard any URL that could not be unwrapped. In general, URLs correspond to online articles, posts, links, or websites.

III. **News domains.** We extract all domain names mentioned inside tweets that correspond to mainstream media sites indexed by Google News [20]. News domains correspond to media outlets, which may be specializing in the coverage of some topics or perspectives.

IV. **YouTube videos.** We extract all URLs mentioned inside tweets that match the pattern `www.youtube.com/watch`. Here, each of these URLs corresponds to a different YouTube video.

The above memes provide different levels of granularity. For example, news domains are very generic, while YouTube videos are fairly specific. In more detail, the set of active users mention 286, 219 unique hashtags, 379, 424 URLs, 18,616 news domains, and 19, 998 YouTube videos. Figure 1A shows the distribution of the number of unique mentions for different types of memes, which follows a power-law distribution. The tail of the distribution, as expected, is the heaviest for news domains, while the lightest for YouTube videos. Moreover, as shown in Figure 1B, the tail of the distribution of the number of followees tweeting at least one of the memes is also a power-law. In the remainder, we consider only such followees.¹ Also, we focus on users whose information network is fairly developed by filtering out any user following less than 20 followees.

Note that, although our methodology does not depend on the particular choice of meme, it does make two key assumptions. First, it assumes we can distinguish whether two memes are equal or differ. Distinguishing certain memes such as hashtags may be trivial but distinguishing others, such as ideas, may be very difficult. Second, it assumes that receiving several copies of the same meme from different users does not provide additional information, even if different users express different opinions about the meme. It would be interesting to relax the second assumption in future work.

Importantly, in 2009 Twitter did not have features such as “Lists” and “Personalized Suggestions”, so the main way users received and processed information was through their feed, for which we have complete data. The drawback of using older data is smaller number of users and social activity.

3. DEFINITIONS OF EFFICIENCY

In this section, we introduce three different notions of efficiency, namely, link, in-flow and delay efficiency. For each type of efficiency, we provide a formal definition and propose a method to approximately compute it with provable guarantees.² Then, we use the methods to investigate the efficiency of Twitter users at acquiring information.

3.1 Link efficiency

Definition: Our definition of link efficiency follows from our prior definition [4]. Consider a user u and the set of unique memes \mathcal{I}_u she is exposed to through her feed in a given time period, by following $|\mathcal{U}_u|$ users. Then, we define the optimal set $\mathcal{U}^l(\mathcal{I}_u)$ as the minimal set of users that, if followed, would expose the user to,

¹Considering all followees leads to qualitatively similar results, but lower absolute values of efficiency.

²Delay efficiency, unlike link and in-flow efficiencies, can be computed exactly.




Figure 2: Our notion of link efficiency, E_u^l . We define link efficiency as $E_u^l = |\mathcal{U}^l(\mathcal{I}_u)|/|\mathcal{U}_u|$, where \mathcal{I}_u is the set of (unique) memes (blue circles) a user u receives in her timeline by following a set of followees $\mathcal{U}_u = \{u_1, \dots, u_5\}$ (left), and $\mathcal{U}^l(\mathcal{I}_u) = \{u_1^*, u_2^*, u_3^*\}$ (right) is the minimal set cover (of users) that, if followed, would provide the same set of memes \mathcal{I}_u . In the illustration, each user u_i posts the memes within the associated ellipsoid. Hence, in this example, the link efficiency value is $E_u^l = 3/5$.

Algorithm 1: Greedy set cover for estimating link efficiency

Input: set of all users \mathcal{U} ; set of unique memes \mathcal{I}_u ; followee set \mathcal{U}_u ; set of memes \mathcal{I}^v posted by user v

Set $\mathcal{U}^l = \emptyset$;

Set $\mathcal{X} = \mathcal{I}_u$;

while $\mathcal{X} \neq \emptyset$ **do**

 Set $v^* = \operatorname{argmin}_{v \in \mathcal{U} \setminus \mathcal{U}^l} \frac{1}{|\mathcal{X} \cap \mathcal{I}^v|}$;

 Set $\mathcal{U}^l = \mathcal{U}^l \cup \{v^*\}$;

 Set $\mathcal{X} = \mathcal{X} \setminus \mathcal{I}_{v^*}$;

end

Output: \mathcal{U}^l

at least, \mathcal{I}_u , and define the link efficiency of a user u at acquiring memes as

$$E_u^l = \frac{|\mathcal{U}^l(\mathcal{I}_u)|}{|\mathcal{U}_u|}, \quad (1)$$

where $0 \leq E_u^l \leq 1$. If the number of users she follows coincides with the number of users in the minimal set, then her efficiency value is $E_u^l = 1$. The larger the original number of followees in comparison with the size of the minimal set, the smaller the link efficiency. Figure 2 illustrates our definition of link efficiency.

Examples of link inefficiency: Our definition captures two types of link inefficiency, which we illustrate by two extreme examples. If a user u follows $|\mathcal{U}_u|$ other users, each of them mentioning different (disjoint) sets of memes, and there is another user $v \notin \mathcal{U}_u$ that cover all the memes the followees cover, then the user’s efficiency will be $E_u^l = 1/|\mathcal{U}_u|$. If a user u follows $|\mathcal{U}_u|$ other users and all these users mention exactly the same memes, then the user’s efficiency will be $E_u^l = 1/|\mathcal{U}_u|$ and $\lim_{|\mathcal{U}_u| \rightarrow \infty} E_u^l = 0$. The former type of link inefficiency is due to following users that individually post too few memes, while the latter is due to following users that collectively produce too many redundant memes.

Computing link efficiency: In practice, computing E_u^l , as defined by Eq. 1, reduces to finding the minimal set of users $\mathcal{U}^l(\mathcal{I}_u)$, which can be cast as the classical minimum set cover problem [16]. Although the minimum set cover problem is NP-hard, we can approximate $\mathcal{U}^l(\mathcal{I}_u)$ using a well-known and efficient greedy algorithm [15], which returns an $O(\log d)$ approximation of the min-




Figure 3: Our notion of in-flow efficiency, E_u^f . We define in-flow efficiency as $E_u^f = f(\mathcal{U}^f(\mathcal{I}_u)) / f(\mathcal{U}_f)$, where \mathcal{I}_u is the set of (unique) memes (blue circles) a user receives in her timeline by following a set of followers $\mathcal{U}_u = \{u_1, \dots, u_5\}$ (left), and $\mathcal{U}^f(\mathcal{I}_u) = \{u_1^*, u_2^*, u_3^*\}$ (right) is the set cover (of users) with the smallest associated in-flow $f(\mathcal{U}^f(\mathcal{I}_u))$ that, if followed, would provide the same set of memes \mathcal{I}_u . In the illustration, each user u_i posts the memes within the associated ellipsoid and the red values in the ellipsoid represent the in-flow of each user. Hence, in this case, the in-flow efficiency value is $E_u^f = 30/60 = 0.5$.

Algorithm 2: Greedy set cover for estimating in-flow efficiency

Input: set of all users \mathcal{U} ; set of unique memes \mathcal{I}_u ; number of tweets N^v posted by user v ; set of memes \mathcal{I}^v posted by user v

Set $\mathcal{U}^f = \emptyset$;
Set $\mathcal{X} = \mathcal{I}_u$;
while $\mathcal{X} \neq \emptyset$ **do**
 Set $v^* = \arg \min_{v \in \mathcal{U} \setminus \mathcal{U}^f} \frac{N^v}{|\mathcal{I}^v \cap \mathcal{X}|}$;
 Set $\mathcal{U}^f = \mathcal{U}^f \cup \{v^*\}$;
 Set $\mathcal{X} = \mathcal{X} \setminus \mathcal{I}_{v^*}$;
end
Output: \mathcal{U}^f

imum size set cover, where $d = \max_{v \in \mathcal{U}} |\mathcal{I}^v|$ is the maximum number of memes posted by any user. Refer to Algorithm 1 for a full description of our procedure to approximate link efficiency.

3.2 In-flow efficiency

Definition: Consider a user u and the set of unique memes \mathcal{I}_u she is exposed to through her feed in a given time period, by following $|\mathcal{U}_u|$ users. Then, we define the optimal set $\mathcal{U}^f(\mathcal{I}_u)$ as the set of users that, if followed, would expose the user to, at least, \mathcal{I}_u , while providing the least amount of tweets per time unit, *i.e.*, the minimum tweet in-flow. In particular, we define the in-flow efficiency of a user u at acquiring memes as

$$E_u^f = \frac{f(\mathcal{U}^f(\mathcal{I}_u))}{f(\mathcal{U}_u)}, \quad (2)$$

where $f(\mathcal{U}_u)$ denotes the amount of tweets produced by users \mathcal{U}_u per time unit (user u 's in-flow) and $0 \leq E_u^f \leq 1$. The in-flow efficiency $E_u^f = 1$ if user u 's in-flow coincides with the amount of tweets per time unit posted by the users in the optimal set $\mathcal{U}^f(\mathcal{I}_u)$. Here, the larger is user u 's in-flow in comparison with the amount of tweets per time unit posted by the users in the optimal set, the lower is her in-flow efficiency. Figure 3 illustrates our definition of in-flow efficiency using an example.




Figure 4: Our notion of delay efficiency, E_u^t . We define delay efficiency as $E_u^t = 1 / (1 + \langle t_i - t_i^0 \rangle_{i \in \mathcal{I}_u})$, where t_i is the time in which a user receives meme i in her timeline, t_i^0 is the time when the meme is first mentioned by a user in the whole social media system, \mathcal{I}_u is the set of (unique) memes (blue circles) a user receives in her timeline by following a set of followers $\mathcal{U}_u = \{u_1, \dots, u_5\}$ (left), and $\mathcal{U}^t(\mathcal{I}_u) = \{u_1^*, \dots, u_3^*\}$ (right) is the set cover (of users) that, if followed, would provide the same set of memes \mathcal{I}_u as early as possible. In the illustration, each user u_i adopts the memes within the associated ellipsoid for the first time after a delay indicated by the red number. Hence, the delay efficiency is $E_u^t = 1 / (1 + 30/13)$.

Examples of in-flow inefficiency: As in the case of link inefficiency, this definition captures several types of in-flow inefficiency. First, it is easy to see that the example of extreme link inefficiency due to following users posting exactly the same memes, also leads to in-flow inefficiency.

Second, there is another type of in-flow inefficiency, which we illustrate by an additional extreme example. Consider user u that follows $|\mathcal{U}_u|$ other users and the amount of tweets produced by these followers has a divergent mean, *e.g.*, it has a Pareto distribution (a power law) with exponent $\alpha \leq 1$. Then, if there exists another set of $|\mathcal{U}_u|$ users mentioning the same unique memes and the amount of tweets produced by them has a non-divergent mean, *e.g.*, it is a Pareto distribution with exponent $\alpha > 1$, the user's efficiency will converge to zero as $|\mathcal{U}_u|$ increases, *i.e.*, $\lim_{|\mathcal{U}_u| \rightarrow \infty} E_u^f = 0$. Asymptotically, an infinite in-flow could be replaced with a finite in-flow that include the same set of unique memes.

Computing in-flow efficiency: In practice, computing the optimal set of users $\mathcal{U}^f(\mathcal{I}_u)$ reduces to solving the weighted set cover problem, which is also NP-hard. Analogously, we can find an approximate solution to $\mathcal{U}^f(\mathcal{I}_u)$ using a greedy algorithm [15], which returns an $O(\log d)$ approximation to the set cover with minimum in-flow, where $d = \max_{v \in \mathcal{U}} |\mathcal{I}^v|$ is the maximum number of memes posted by any user. Refer to Algorithm 2 for a full description of our procedure to approximate in-flow efficiency with an approximation factor $O(\log d)$.

3.3 Delay efficiency

Definition: Consider a user u and the set of unique memes \mathcal{I}_u she is exposed to through her feed in a given time period, by following $|\mathcal{U}_u|$ users. Then, we define the optimal set $\mathcal{U}^t(\mathcal{I}_u)$ as the set of users that, if followed, would expose the user to, at least, \mathcal{I}_u , with the smallest time delay. Here, we define the delay at acquiring a meme provided by a set \mathcal{U}_u as the difference between the time when a user in \mathcal{U}_u first mentions the meme and the time when the meme was first mentioned during the given time period by any user in the whole social media system. We then define the delay

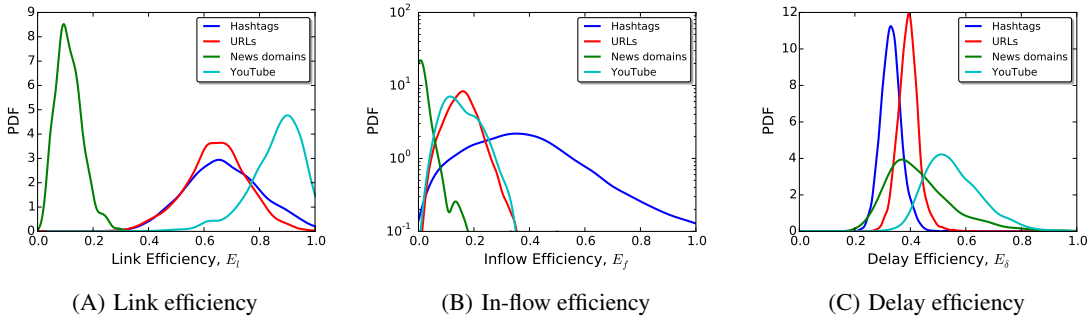


Figure 5: The distributions of link, in-flow, and delay efficiencies for the four types of memes.




Figure 6: The distribution of the ratio between the number of received tweets and unique memes.

efficiency of a user u at acquiring memes as

$$E_u^t = \frac{1}{1 + \langle t_i - t_i^0 \rangle_{i \in \mathcal{I}_u}}, \quad (3)$$

where t_i is the time a user in \mathcal{U}_u first mentions meme i , t_i^0 is the time when the meme is first mentioned by a user in the whole social media system, and $\langle t_i - t_i^0 \rangle_{i \in \mathcal{I}_u}$ is an average delay over all memes received by user u , measured in days. The delay efficiency $E_u^t = 1$ if the followees of the user u are the first to post the set of memes \mathcal{I}_u in the whole system. The delay efficiency becomes lower than 1 when the user is exposed to the memes at later times than their time of birth. The larger is the average delay of received memes, the smaller is the delay efficiency. Figure 4 illustrates our definition of delay efficiency using an example.

Computing delay efficiency: In this case, we can compute the delay efficiency directly by finding when each of the memes appeared for the first time in the system, without resorting to approximation algorithms, as in the case of link and in-flow efficiencies. One can query the first time of appearance for each meme in $O(1)$ by building a mapping between memes and their first time of appearance in a hashtable.

3.4 Efficiency of Twitter users

Once we have three definitions of users' efficiency, defined by Eqs. 1-3, we use them to investigate how efficient Twitter users are at acquiring four different types of memes: hashtags, URLs, news domains and YouTube videos.

First, we estimate the empirical probability density function³ (PDF) for each type of efficiency and meme. We show the results in Figure 5, in which we find several interesting patterns. First, all PDFs resemble a normal distribution, however, their peaks (modes)

³The PDFs have been empirically estimated using kernel density estimation [7].




Figure 7: The average link and in-flow efficiencies versus the percentage of covered memes.

and widths (standard deviations) differ across efficiencies and type of memes. For most users and most types of memes, the efficiency value is well below one, giving empirical evidence that users are typically sub-optimal. Second, while the PDFs for link (Figure 5A) and delay (Figure 5C) efficiencies look quite similar, the in-flow efficiency differs significantly (Figure 5B). Third, users are most efficient at acquiring YouTube videos, followed by URLs and hashtags, and news domains. This order coincides with the ordering of the exponents of the corresponding power-law distribution of memes' popularity (Figure 1A), *i.e.*, the exponent of the power-law (its absolute value) is the highest for YouTube links, followed by URLs and hashtags, and finally for news domains. Note that the higher the exponent is, the higher the proportion of non-popular memes with respect to the popular ones, and thus one can conclude that users are more efficient at acquiring non-popular memes than popular memes. A plausible explanation is that users posting non-popular memes are likely to be included in the optimal set, since there is nobody else who posts these memes and, as a consequence, the optimal set differs less from the original set of followees. Moreover, notice that in Figure 5B, the in-flow efficiency of both news domains and YouTube videos is shifted to the left (*i.e.*, presents much lower efficiency values) compared to the link efficiency in Figure 5A. This shift is due to the fact that, as shown in Figure 6, ratio between the total number of received tweets and unique memes is much larger for unique news domains and YouTube video memes than for hashtags and URLs, which in turn, translates into a lower in-flow efficiency.

In the above measurements, we estimated the probability density functions of user's efficiency considering full coverage of the received memes. Importantly, it is straightforward to extend our definitions of link and in-flow efficiencies to account for partial coverage, by simply considering a set of users that, if followed, would

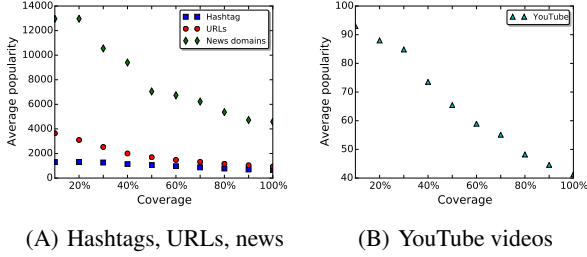


Figure 8: The average popularity of covered memes as a function of the percentage of covered memes.

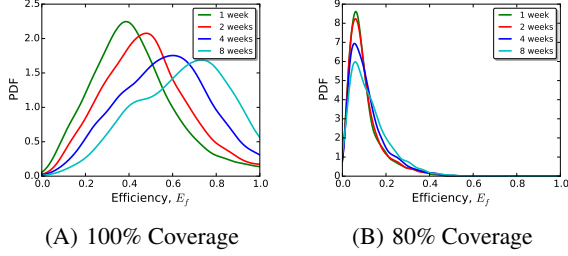


Figure 9: In-flow efficiency measured for hashtags appearing in the time periods of different lengths. The results for other efficiencies and meme types are qualitatively the same.

expose the user to, at least, a percentage of the unique memes \mathcal{I}_u , by stopping the greedy algorithm whenever the given percentage is reached. Note, however, that computing the efficiency for a partial coverage based on the full set of followers would be unfair, since some of the followers in \mathcal{U}_u may be not tweeting any of the covered memes. Thus, for the purpose of computing the efficiencies for partial coverage, we take into account only the users in \mathcal{U}_u who tweet at least one of the covered memes. Figure 7 shows the average link and in-flow efficiencies against coverage for the same four memes. As one may have expected, the higher the coverage, the higher the link and in-flow efficiency, since the memes that are covered first by the greedy algorithm are the popular ones, as shown in Figure 8. This result confirms that users are more efficient at acquiring less popular information, but less so at acquiring more popular information. A plausible explanation is that less popular information is produced by only a handful of users and so the optimization is limited to this set of users.

Finally, we investigate if our results are consistent across different time periods. In particular, we measure user efficiency based on time periods of different lengths: one, two, four and eight weeks. In Figure 9A, we can observe that as we increase the period, there are more unique memes there to cover, which results in an increase in the efficiency. However, the distribution of efficiency is nearly unchanged for 80% coverage (Figure 9B). Thus, the findings presented in this study are qualitatively robust to the choice of time period. Additionally, we find that the choice of the week does not influence the distributions of efficiency, however, these results are not shown due to space limitation.

4. CROSS-EFFICIENCY

In our definitions of efficiency, the optimal set for a given user is the set of users that minimizes the number of links, in-flow or delay, while covering the same set of unique memes. However,

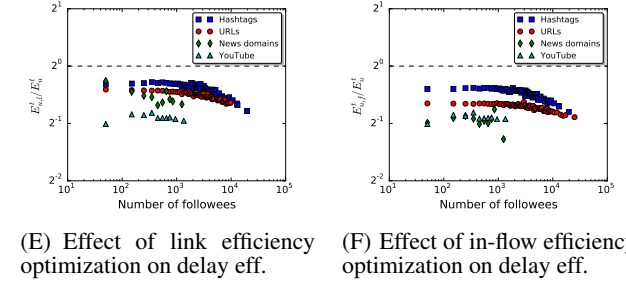
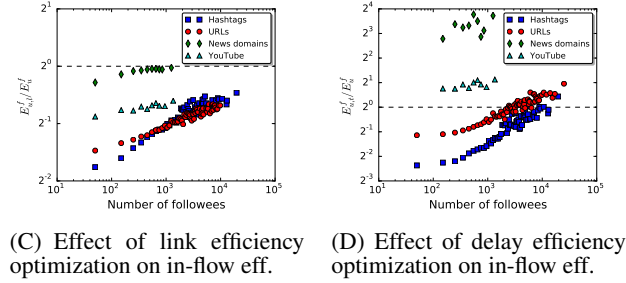
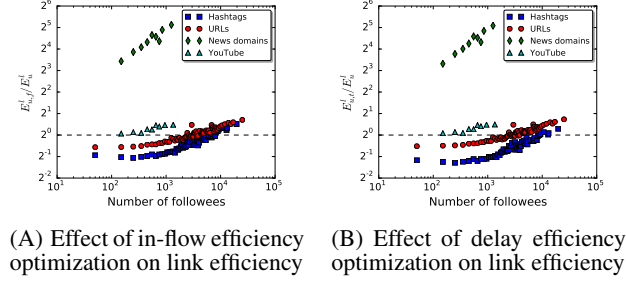


Figure 10: The effect of optimization of one of the efficiencies on another efficiency, plotted as the ratio of the efficiency in the optimized network and the original network against the number of followers. The dashed line marks the ratio equal to 1, which corresponds to the lack of change in the efficiency due to the respective optimization.

this naturally raises the question as to how efficient the optimal sets for a given definition of efficiency are in terms of the other definitions. For example, how efficient is the link-optimal set with respect to in-flow or delay efficiency? In this section, we first address this question, then introduce the idea of finding sets of users that jointly optimize multiple notions of efficiency, and finally develop a heuristic algorithm that simultaneously improves both in-flow and delay efficiency of users.

4.1 Cross-efficiency of optimal sets

Given a user u and the set of unique memes \mathcal{I}_u she is exposed to in a given time period, our definitions of efficiency compare the original set of followers with the optimal sets $\mathcal{U}^l(\mathcal{I}_u)$, $\mathcal{U}^f(\mathcal{I}_u)$ and $\mathcal{U}^d(\mathcal{I}_u)$ in terms of number of links, in-flow and average delay, respectively. Here, we assess the efficiency of the optimal sets for each definition of efficiency in terms of the other definitions, which we call *cross-efficiencies*. More specifically, we compute the link efficiency of the optimal sets for in-flow and delay efficiency, *i.e.*,

$$E_{u,f}^l = \frac{|\mathcal{U}^l(\mathcal{I}_u)|}{|\mathcal{U}^f(\mathcal{I}_u)|} \quad \text{and} \quad E_{u,t}^l = \frac{|\mathcal{U}^l(\mathcal{I}_u)|}{|\mathcal{U}^d(\mathcal{I}_u)|},$$

the in-flow efficiency of the optimal sets for link and delay efficiency, *i.e.*,

$$E_{u,l}^f = \frac{f(\mathcal{U}^f(\mathcal{I}_u))}{f(\mathcal{U}^l(\mathcal{I}_u))} \quad \text{and} \quad E_{u,t}^f = \frac{f(\mathcal{U}^f(\mathcal{I}_u))}{f(\mathcal{U}^t(\mathcal{I}_u))},$$

and the delay efficiency of the optimal sets for link and in-flow efficiency, *i.e.*,

$$E_{u,l}^t = \frac{1}{1 + \langle t_i^l - t_i^0 \rangle_{i \in \mathcal{I}_u}} \quad \text{and} \quad E_{u,t}^t = \frac{1}{1 + \langle t_i^t - t_i^0 \rangle_{i \in \mathcal{I}_u}},$$

where t_i^l is the time a user in $\mathcal{U}^l(\mathcal{I}_u)$ first mentions meme i and t_i^t is the time a user in $\mathcal{U}^t(\mathcal{I}_u)$ first mentions meme i .

Typically, we would like to know if an efficiency of an optimized information network is increased in comparison with the efficiency of the original network. Thus, in the remainder of this section we focus on measuring the ratio of an efficiency of the optimized and original networks. If the ratio is higher than one for the given optimization algorithm, then the corresponding efficiency is improved by that algorithm with respect to the original set of followees. If the ratio is below one then the respective efficiency is decreased by the optimization algorithm.

We measure the ratio between the efficiency of the optimal sets and the original set of followees for the three definitions of efficiency (see Figure 10). The ratio tends to be below one or close to one for most meme types, which indicates that optimizing for one definition of efficiency generally results in decreased efficiency with respect to the other two definitions.

However, there are a few exceptions. For example, in terms of link and in-flow efficiency, the optimal sets for news domains are more efficient than the original sets (diamonds in Figures 10A, 10B, and 10D). This observation may happen due to the following reason: news domains tend to be more popular than the other types of memes (as shown in Figure 1). As a consequence, a user may receive multiple copies of the same news domain from various followees, and it is very easy to find efficient sets in terms of in-flow; it is enough to simply remove some of their followees from the network to improve both link and in-flow efficiencies.⁴

Finally, in Figures 10A-10D, we note that the improvement in the link and in-flow efficiencies tends to grow with the number of followees due to the increased number of redundant (non-unique) information received by the users who follow many other people. However, in Figures 10E-10F, the improvement in the delay efficiency tends to drop with the number of followees because it is likely that users who receive many copies of the same meme receive it early on. Thus, for users with many followees, it is harder to improve the delay efficiency.

4.2 Joint-optimization of efficiencies

So far, we have looked for optimal sets of users in terms of a single efficiency (be it link, in-flow or delay). Moreover, in the previous section, we have shown that optimal sets in terms of a single efficiency typically decrease the other efficiencies. Therefore, one could imagine developing an algorithm a looking for sets of users that are optimized with respect to several efficiencies; in other words, a multi-objective algorithm. Given such an algorithm a , we could compute the efficiency of the optimal set $\mathcal{U}_a(\mathcal{I}_u)$ with

⁴In fact, over 85% of users receive less unique news domains than they have followees.

Algorithm 3: Greedy set cover for optimizing in-flow and delay efficiencies

Input: set of all users \mathcal{U} ; set of unique memes \mathcal{I}_u ; set of memes \mathcal{I}^v posted by user v

Set $\mathcal{U}^* = \emptyset$;

Set $\mathcal{X} = \mathcal{I}_u$;

while $\mathcal{X} \neq \emptyset$ **do**

Set $v^* = \arg \min_{v \in \mathcal{U} \setminus \mathcal{U}^f} \frac{N_v^\alpha T_v^\beta}{|\mathcal{I}^v \cap \mathcal{X}|}$;

Set $\mathcal{U}^* = \mathcal{U}^* \cup \{v^*\}$;

Set $\mathcal{X} = \mathcal{X} \setminus \mathcal{I}_{v^*}$

end

Output: \mathcal{U}^*

respect to single quantities, *i.e.*,

$$E_{u,a}^l = \frac{|\mathcal{U}^l(\mathcal{I}_u)|}{|\mathcal{U}^a(\mathcal{I}_u)|}, \quad E_{u,a}^f = \frac{f(\mathcal{U}^f(\mathcal{I}_u))}{f(\mathcal{U}^a(\mathcal{I}_u))} \quad \text{and}$$

$$E_{u,a}^t = \frac{1}{1 + \langle t_i^a - t_i^0 \rangle_{i \in \mathcal{I}_u}}$$

where t_i^a is the time a user in $\mathcal{U}^a(\mathcal{I}_u)$ first mentions meme i . Ideally, we would like to find optimal sets that are efficient with respect to the considered quantities. Here, as a proof of concept, we next develop a heuristic method to find sets of users optimized with respect to both in-flow and delay.

Joint optimization of in-flow and delay efficiency

We leverage the greedy algorithm from the weighted set cover problem to design a heuristic method that finds sets of users with high in-flow and delay efficiencies, while delivering the same unique memes to the user (refer to Algorithm 3). In particular, in the heuristic method, the weights are powers of tweets in-flow N_v^α and average delay T_v^β over all unique memes produced by the user v . The exponents α and β can be readily adjusted to induce higher or lower in-flow efficiency and delay efficiency, respectively. Here, we experiment with $\alpha = 1$ and $\beta = 0.5$, which achieves a good balance between in-flow and delay efficiency.

We summarize the ratio of link, in-flow and delay efficiency of the set of users provided by our heuristic method and the original set of followees in Figure 11. We discuss several interesting observations. First, since the algorithm does not optimize with respect to the number of links, the link efficiency is not improved by this algorithm, *i.e.*, the ratio between the link efficiency of the set provided by the heuristic method and the link efficiency of the original set of followees ratio is around or below 1 for three out of four meme types. Second, we find that both in-flow and delay efficiencies are significantly increased over the efficiency of the original set of followees for all types of memes. The in-flow efficiency is, on average, 7.4-times higher for news domains, 1.8-times higher for hashtags, 1.3-times higher for URLs, 1.2-times higher for YouTube videos. The delay efficiency is, on average, 1.8-times higher for news domains, 1.4-times higher for hashtags, 1.4-times higher for URLs, 1.2-times higher for YouTube videos. There is always an increase in in-flow and delay efficiencies independently of the number of followees that the users have originally. However, while the improvement in the in-flow efficiency tends to be larger for users with many followees, the improvement in the delay efficiency is larger for users with fewer followees. Thus, we conclude that our algorithm increases both in-flow and delay efficiency of users.

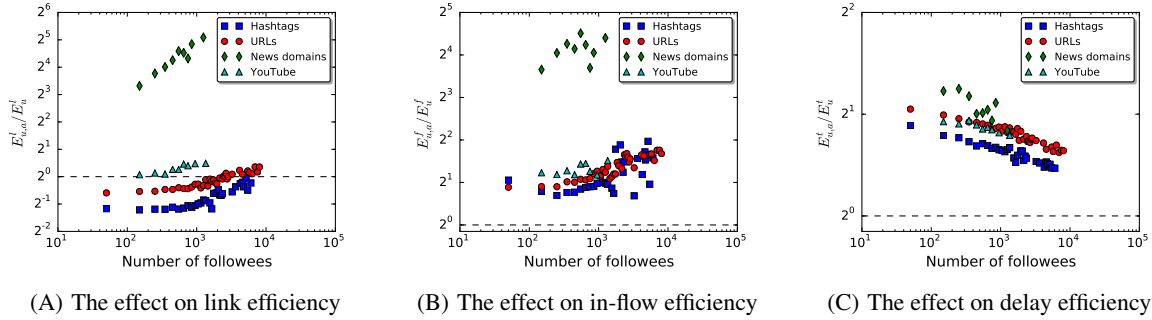


Figure 11: The effect of optimization of both in-flow and delay efficiencies on different types of efficiencies, plotted as the ratio of the affected efficiency of the optimized network and the original network against the number followers.




Figure 12: Ego-networks for a Twitter user (red node). Some users (black nodes) only belong to one of the ego-networks while others (gray nodes) belong to both. The original ego-network contains more triangles than the ego-network induced by the minimal set cover, whose structure is closer to a star.

5. STRUCTURE OF EGO NETWORKS: ORIGINAL VS. OPTIMIZED

In the previous sections, we have introduced three meaningful definitions of efficiency and applied them to show that Twitter users tend to choose their information sources inefficiently. In this section, we investigate the rationale behind this sub-optimal behavior by comparing the structure of the user’s ego-networks associated with both the original set of followers and the sets optimized for efficiency. Here, we define a user’s ego-network as the network of connections (who-follows-whom) between the ego user and her followers.

First, as an example, we take one particular user and illustrate the structure of her original ego-network and the ego-network of an optimal set in terms of link efficiency (Figure 12). By visual comparison of both ego-networks, we can see that while the ego-network induced by the optimal set displays a structure much closer to a star, the original ego-network contains many more triangles and higher clustering coefficient. Due to its proportionally lower number of triangles, the optimal set is not *discoverable* by triadic closure [26, 13, 24] or information diffusion [5], which have been recently shown to be two major driving forces for link creation in social networks [28, 23, 3].

Remarkably, this phenomenon happens systematically across all users, efficiency definitions, and types of memes, as displayed in Figure 13, which shows the distribution of local clustering coefficient (LCC) for the users’ original ego-networks and the ego-networks induced by different optimized sets. We find that while

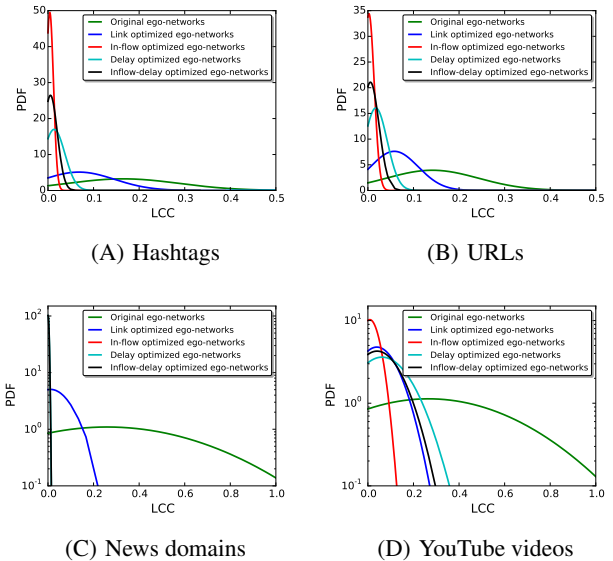


Figure 13: The distributions of local clustering coefficient (LCC) of the original ego-network (green circles) and the ego-networks optimized for link (blue circles), in-flow (red circles), delay (teal circles), and inflow-delay efficiency (black circles) for: (A) hashtags, (B) URLs, (C) news domains, (D) and YouTube videos.

the LCC distribution for the original ego-networks is well spread and centered at 0.15 – 0.30, the LCC distributions for the ego-networks induced by the optimal sets are skewed towards zero.⁵ One could still think that this is simply a consequence of differences in the number of followers, *i.e.*, the size of the ego-network. However, Figure 14 rules out this possibility by showing a striking difference of several orders of magnitude between the LCC of the original ego-networks and the ego-networks induced by the optimal sets across a wide range of number of followers. These findings suggest that the way in which social media users discover new people to follow (*e.g.*, triadic closure or information diffusion) or receive recommendations (*e.g.*, pick people in a 2-hop neighbor-

⁵Note that the distribution of clustering coefficient of inflow-delay optimized network is located between distributions of in-flow optimized and delay optimized ego-networks.

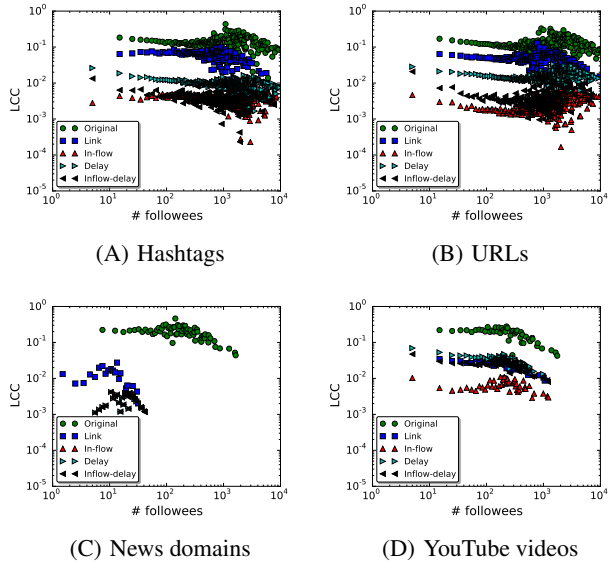


Figure 14: Average local clustering coefficient versus the number of followees in the original ego-network (green circles) and the ego-networks optimized for link (blue squares), in-flow (red triangles), delay (teal triangles), and inflow-delay efficiency (black triangles).

hood [1]) can lead to sub-optimal information networks in terms of (link, in-flow and delay) efficiency.

We have argued that optimal sets typically differ from the original set of followees due to their low number of triangles in the associated ego-networks, and thus lack of discoverability. However, are optimal sets with higher number of triangles in efficient ego-networks easier to discover for users? Figure 15 answers this question positively by showing the average local clustering coefficient in the ego-network induced by the optimal set against the overlap between the users in the optimal set and the original set of followees. Here, by overlap we mean the fraction of users in the optimal set that are also in the original set of followees. In particular, we find a positive correlation (Pearson’s $0.07 < r < 0.55$, $p < 10^{-10}$) between the local clustering coefficient and the overlap, which indicates that if the nodes in the optimal set are *discoverable* through triadic closure, the user may be more likely to find them and decide to follow them.

6. DISCUSSION

We have defined three intuitive notions of user’s efficiency in social media – link, in-flow and delay efficiency – to assess how *good* users are at selecting who to follow within the social media system to acquire information. Our framework is general and applicable to any social media system where every user *follows* others within the system to receive the information they produce. We have then leveraged our notions of efficiency to help us in understanding the relationship between different factors, such as the popularity of received information and the users’ ego-networks structure.

Here, we have focused on three definitions of efficiency (link, in-flow, and delay). However, we could leverage this idea to define more complex notions of efficiency. For example, we could define efficiency in terms of diversity, *i.e.*, it would be interesting to find the set of users that, if followed, would cover the same unique memes while maximizing the diversity of topics or per-

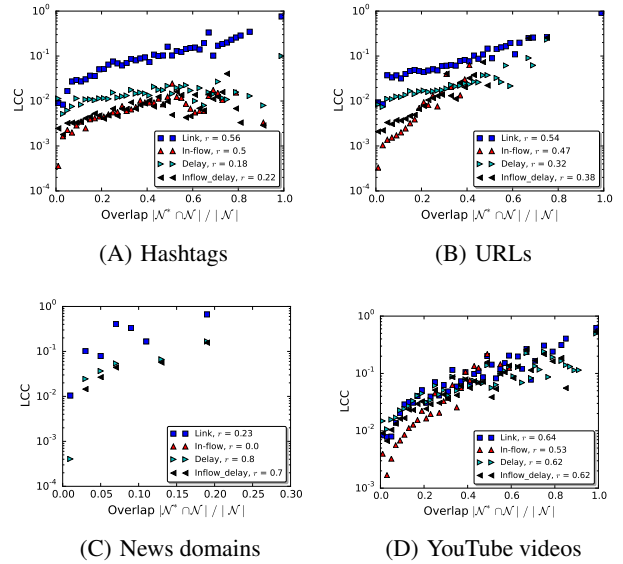


Figure 15: Local clustering coefficient of the optimized ego-networks versus the overlap between the original ego-network and the ego-networks optimized for link (blue squares), in-flow (red triangles), delay (teal triangles), and inflow-delay efficiency (black triangles).

spectives that are delivered with the memes, and then compare this set with the original set of followees in terms of diversity. This would provide a framework to mitigate the effects of the filtering bubble and echo chamber present in current social media systems. Moreover, some of the memes could be treated preferentially over other memes. This could be achieved by means of covering a list of non-unique memes favoring repetitions of a preferential subset of memes, *e.g.*, memes matching the user’s interests should be delivered to the user more often. Remarkably, these more complex notions of efficiency can often be expressed as integer linear programs, similarly to the minimal set cover problem, which can be solved using relaxation methods with provable guarantees [27].

Additionally, we have introduced a heuristic method that improves both in-flow and delay efficiency of users, while still delivering them the same unique memes. Similar heuristics can be naturally designed to optimize efficiency with respect to multiple quantities (be it link, in-flow, delay, or diversity). In this context, it would be very interesting to design methods with provable guarantees to find sets of users that are optimal with respect to multiple quantities.

Our work also opens other interesting venues for future work. For example, we have defined and computed a measure of efficiency for each user independently. However, one could also think on global notions of efficiency for the Twitter information network as a whole, perhaps using a multi set cover approach. Finally, since we have applied our framework to study information efficiency only on Twitter, it would be interesting to study information efficiency of other microblogging services (Weibo, Pinterest, Tumblr) and social networking sites (Facebook, Google+).

7. REFERENCES

- [1] L. A. Adamic and E. Adar. Friends and neighbors on the web. *Social networks*, 25(3):211–230, 2003.
- [2] E. Agichtein, C. Castillo, D. Donato, A. Gionis, and

- G. Mishne. Finding high-quality content in social media. In *Proceedings of the 1st International Conference on Web Search and Data Mining*, pages 183–194, 2008.
- [3] D. Antoniadis and C. Drovolis. Co-evolutionary dynamics in social networks: A case study of twitter. *arXiv preprint arXiv:1309.6001*, 2013.
- [4] M. Babaei, P. Grabowicz, I. Valera, and M. Gomez-Rodriguez. On the users’ efficiency in the twitter information network. In *Proceedings of the 9th International AAAI Conference on Weblogs and Social Media*, 2015.
- [5] E. Bakshy, I. Rosenn, C. Marlow, and L. Adamic. The role of social networks in information diffusion. In *Proceedings of the 21st international conference on World Wide Web*, pages 519–528, 2012.
- [6] R. Bosagh Zadeh, A. Goel, K. Munagala, and A. Sharma. On the precision of social and information networks. In *Proceedings of the first ACM conference on Online social networks*, pages 63–74. ACM, 2013.
- [7] A. W. Bowman and A. Azzalini. *Applied smoothing techniques for data analysis*. Clarendon Press, 2004.
- [8] C. Castillo, M. Mendoza, and B. Poblete. Information credibility on twitter. In *Proceedings of the 20th International Conference on World Wide Web*, pages 675–684, 2011.
- [9] M. Cha, H. Haddadi, F. Benevenuto, and P. K. Gummadi. Measuring User Influence in Twitter: The Million Follower Fallacy. In *Proceedings of the 4th International AAAI Conference on Weblogs and Social Media*, pages 10–17, 2010.
- [10] N. A. Christakis and J. H. Fowler. Social network sensors for early detection of contagious outbreaks. *PloS one*, 5(9):e12948, 2010.
- [11] M. Farajtabar, M. Gomez-Rodriguez, N. Du, M. Zamani, H. Zha, and L. Song. Back to the past: Source identification in diffusion networks from partially observed cascades. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics*, 2015.
- [12] M. Gomez-Rodriguez, K. Gummadi, and B. Schoelkopf. Quantifying Information Overload in Social Media and its Impact on Social Contagions. In *Proceedings of the 8th International AAAI Conference on Weblogs and Social Media*, pages 170–179, 2014.
- [13] M. S. Granovetter. The strength of weak ties. *American journal of sociology*, pages 1360–1380, 1973.
- [14] N. Hodas and K. Lerman. How visibility and divided attention constrain social contagion. In *Proceedings of the 2012 ASE/IEEE International Conference on Social Computing*, pages 249–257, 2012.
- [15] D. S. Johnson. Approximation algorithms for combinatorial problems. In *Proceedings of the fifth annual ACM symposium on Theory of computing*, pages 38–49. ACM, 1973.
- [16] R. M. Karp. *Reducibility among combinatorial problems*. Springer, 1972.
- [17] M. Kearns. Experiments in social computation. *Communications of the ACM*, 55(10):56–67, 2012.
- [18] H. Kwak, C. Lee, H. Park, and S. Moon. What is twitter, a social network or a news media? In *Proceedings of the 19th International Conference on World Wide Web*, pages 591–600, 2010.
- [19] K. Lerman and T. Hogg. Leveraging position bias to improve peer recommendation. *PLoS One*, 9(6), 2014.
- [20] J. Leskovec, L. Backstrom, and J. Kleinberg. Meme-tracking and the dynamics of the news cycle. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 497–506, 2009.
- [21] J. Leskovec, A. Krause, C. Guestrin, C. Faloutsos, J. VanBriesen, and N. Glance. Cost-effective outbreak detection in networks. In *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 420–429. ACM, 2007.
- [22] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee. Measurement and analysis of online social networks. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, pages 29–42, 2007.
- [23] S. A. Myers and J. Leskovec. The bursty dynamics of the twitter information network. In *Proceedings of the 23rd International Conference on World Wide Web*, pages 913–924, 2014.
- [24] D. M. Romero and J. Kleinberg. The directed closure process in hybrid social-information networks, with an analysis of link formation on twitter. In *Proceedings of the 4th International AAAI Conference on Weblogs and Social Media*, 2010.
- [25] D. M. Romero, B. Meeder, and J. Kleinberg. Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter. In *Proceedings of the 20th International Conference on World Wide Web*, pages 695–704, 2011.
- [26] G. Simmel. *The Sociology of Georg Simmel*. Free Press of Glencoe, 1950.
- [27] V. V. Vazirani. *Approximation algorithms*. Springer, 2001.
- [28] L. Weng, J. Ratkiewicz, N. Perra, B. Gonçalves, C. Castillo, F. Bonchi, R. Schifanella, F. Menczer, and A. Flammini. The role of information diffusion in the evolution of social networks. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 356–364, 2013.
- [29] R. West and J. Leskovec. Automatic versus human navigation in information networks. In *Proceedings of the 6th International AAAI Conference on Weblogs and Social Media*, 2012.
- [30] R. West and J. Leskovec. Human wayfinding in information networks. In *Proceedings of the 21st International Conference on World Wide Web*, pages 619–628, 2012.