

Article

Interactional and Informational Attention on Twitter

Agathe Baltzer ^{1,2,†}, Márton Karsai ^{1,*,†}  and Camille Roth ^{3,*,†} 

¹ Univ Lyon, Inria, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR 5668, F-69007 Lyon, France

² Mines ParisTech, 75006 Paris, France

³ CNRS and Computational Social Science Team, Centre Marc Bloch, UMIFRE CNRS-MAEE 14, Friedrichstrasse 191, D-10117 Berlin, Germany

* Correspondence: marton.karsai@ens-lyon.fr (M.K.); roth@cmb.hu-berlin.de (C.R.)

† These authors contributed equally to this work.

Received: 24 May 2019; Accepted: 24 July 2019; Published: 29 July 2019



Abstract: Twitter may be considered to be a decentralized social information processing platform whose users constantly receive their followees' information feeds, which they may in turn dispatch to their followers. This decentralization is not devoid of hierarchy and heterogeneity, both in terms of activity and attention. In particular, we appraise the distribution of attention at the collective and individual level, which exhibits the existence of attentional constraints and focus effects. We observe that most users usually concentrate their attention on a limited core of peers and topics, and discuss the relationship between interactional and informational attention processes—all of which, we suggest, may be useful to refine influence models by enabling the consideration of differential attention likelihood depending on users, their activity levels, and peers' positions.

Keywords: attention; influence; ego-centered networks; twitter study; information spreading

1. Introduction

Attention and activity seem to obey possibly conflicting dynamics. On the one hand, a wealth of results have accumulated to suggest that social attention is generally bounded, whereby humans are not able to devote their active time to more than a certain number of peers and topics [1–5], pointing at attention as a zero-sum game constrained by limited temporal resources. On the other hand, several studies hint at reinforcement mechanisms between various attentional and activity channels, especially when comparing online with offline sociability (or social capital), which appear to be correlated [6–9]: in this regard, attention and/or activity may breed more attention and/or activity.

This paper aims at shedding further light on human attentional patterns by encompassing both interactional and informational aspects, and more precisely by describing the possible existence of conflicting vs. reinforcing cognitive constraints both in social and topical terms. To this end, we focus on a popular, observable online platform, Twitter, which features both social network and publication capabilities, thus making it possible to discuss attention from the joint perspective of interaction and information processing. Twitter allows users to constitute their own personal set of sources of which they want to follow the publications. They can also potentially dispatch these publications, thereby indicating that they indeed paid specific attention to specific users. This will enable us to make a key distinction between potential and actual attentional patterns. By doing so, we more broadly aim at describing how diverse users are in distributing their attentional resources to peers and topics: within certain cognitive limitations, such a platform also exhibits heterogeneous distributions of roles and attentional patterns.

The next section will be devoted to a brief review of the relevant state of the art on this matter. Section 3 will describe the empirical data and the main definitions we use. Section 4 focuses on social attention, while Section 5 introduces the notion of semantic attention and discusses its correlation with its social counterpart, which contributes to more generally address the above debates and evoke further research (Section 7).

2. Related Work

Size of social interaction networks. A sizable literature shows that the number of connections in human ego-centered social interaction networks spans over several orders of magnitude, be it for scientific collaborations [10], e-mail interlocutors [11], content sharing platforms [12], online social networks [13], among others. These studies have principally focused on the number of individuals that one may have known or have interacted with, showing that there is wide variation in human interaction *potential*. When focusing on *actual* interactions and, more precisely, on actively sustained interactions, a diverse array of studies nonetheless converges on the conception that the active core of an individual's interaction network is of bounded size. This follows from Hill & Dunbar's seminal study on the exchange of Christmas cards [1], which shows that individuals may actively devote their actual attention to only a portion of their potential acquaintances. Upper bounds on active connections are generally thought to be in the vicinity of a hundred people, even if this number greatly depends on various features such as, obviously, network type, connection strength thresholds and socio-demographic features, including age and gender. On the lower bound it may go to a dozen or even a handful of contacts when focusing on the most inner layers e.g., where financial support may be sought [2]. In terms of non-face-to-face communication, human interaction capacity is shown to be similarly bounded, be it when examining phone call records [4] or online social network friends [3]. In other words, a variety of constraints, be they physical or cognitive, may contribute to the more or less acute reduction that an individual's social network undergoes when going from potential to actually sustained connections.

Attention dynamics on Twitter. The issue of interactional constraints has been addressed over the recent years on various online platforms, where Twitter has increasingly served as a prototype of observable online networking processes. This scholarship has taken place within a broader questioning on attention dynamics and its possible bounds. On Twitter, social attention is subjected to a threshold similar to what is found in other contexts: for example, reciprocated (conversational) links tend to plateau after a threshold of a couple of hundreds of connections [14]. Semantic attention exhibits regularities at the collective level that may be interpreted as the result of underlying synchronization forces, for instance in terms of typical aggregate sigmoidal patterns of growth and decrease of the global occurrence of some topic [15] and its burstiness [16]. However, a broader picture of the constraints that apply *individually* to attention in this context is only partially known. The attention devoted by Facebook users to their friends tends to follow a power-law distribution, i.e., there is a geometrically decreasing interest as one goes down the list of a user's neighbors [17], while a core of about a dozen of top users appears to consistently gather a significant portion of that attention. Similar features may be found in cell phone communication patterns [18]. In a more open digital public space such as Twitter, the number of unique users one pays attention to varies significantly across the platform as well as during specific events [19]. Yet, little is nonetheless known on the way Twitter users distribute and, possibly, bound their weighted attention among the sources they follow. On the semantic side too, the existence of individual limitations has not really been addressed per se, even if there are detailed accounts of the diversity of topic use in online platforms at large [20,21] and specifically on Twitter [5]. The potential combination of constraints on both the social and semantic sides remains also generally unexplored.

Influence studies. The modeling of social contagion is often configured as a process where the number of neighbors of a given individual matters in a uniform way, whereby all alters have an identical potential impact on ego—be it in the canonical threshold or cascade models [22] or in the so-called “complex contagion” models [23]. In other words, whereas temporal patterns (e.g., repetition, burstiness) and topology (e.g., centrality, clustering) do generally matter, attention is rather considered equal across peers. Attentional constraints have only recently been studied from a diffusion perspective, by examining how limits on individual processing capacities may affect the rate of propagation of information from a user to the other. For instance, [24] shows that there is a decreasing probability of retweeting for users who follow a larger number of people and attempt at estimating response times from observed retweets, while [25] introduces a queuing process aimed at reflecting the sequential processing of information carried out by users in an environment when confronted with an accumulation of information in-flows. While some recent social contagion models specifically introduced the possibility of heterogeneous inter-personal influence [26], corresponding empirical characterizations appear to be still missing. On the whole, existing empirical studies shed an important light on the ego-centered forwarding of information at a short time-scale, yet they do not take into account a finer understanding of how users may heterogeneously distribute their attention across their neighbors or across topics. We intend here to contribute to this question as well.

3. Definitions and Empirical Protocol

3.1. Dataset

We analyze in this study a large corpus stemming from Twitter. It is a popular online news and social networking platform, which enables various types of users to create an account and briefly describe themselves, to publish messages, or “tweets”, with restricted length. Tweets may feature additional content (such as URLs, videos, photos), to subscribe to content generated by other users and to interact with them in various manners (republishing their posts, mentioning them, initiating conversations). Our dataset has been collected over two months, between $t = 6$ June 2016 and $t' = 7$ August 2016 via the Twitter PowerTrack API provided by DataSift with an access rate of about 15%. We focused on users who have self-declared in their account settings that they are located in the time zones GMT and GMT+1. Please note that this is not equivalent to the geolocalization of single tweets, neither to the self-reported default location of users. While such precise locational information is optional to share, time zones are requested to set for each user to match the time of their posting activity. We additionally focused on tweets written in French. The dataset contains 8,233,354 tweets and 8,698,610 retweets.

3.2. Definitions and Notations

Focusing on retweets as attentional markers. We principally focus on content flows exhibited by publication and republication dynamics, i.e., attention and influence dynamics between users. This is admittedly a proxy of a broader notion of attention. Indeed, studying attention naturally requires to define a perimeter of observation: studies using Christmas cards overlook face-to-face, e-mail or phone interactions, and vice versa. Similarly, even when focusing on just Twitter and its quite simple interaction grammar, many distinct attention channels may be considered, and thus many signals of how users pay attention. Users may read posts, their linked content (articles, videos, etc.), they may converse with other users, read posts published by users they do not follow or read the notifications they receive when they are being mentioned in other tweets, and so on and so forth. Without a comprehensive monitoring protocol able to track Twitter users and, especially, both the content they are exposed to and their actual behavior (while factoring in their varying levels of involvement on the platform: from casual to very active users)

we are bound to use secondary signals. From an attentional perspective, we thus decided to focus on user posts as *visible* traces of user-centric activity. We further focus on retweets, for two main reasons:

- We are interested in the cognitive filtering process that occurs between followed sources (and followees' publications) and the actual attention devoted to them. We contend that this constitutes a consistent system that enables us to properly compare what users are exposed to with what they retain. Retweeting is admittedly an ambiguous activity: it has long been considered to be influenced by a variety of temporal and individual factors, either observed [27,28] or hypothesized [29], and has been shown to range from simple acknowledgement to tentative conversation engagement [30]. Yet, it also positively denotes the fact that someone tangibly read a tweet (not necessarily the linked content) among the sources they follow and is minimally interested in the topics evoked in that tweet.
- We jointly consider interactional and informational attention. In this respect, focusing on retweets provides a uniform way to discuss social and semantic attention. In the case of semantic attention, we will nonetheless later show that results are consistent when considering all tweeting activities or just retweets: this further suggests that it remains sound to study both types of attention through retweets only.

Follower and retweet networks. To this end, we introduce two key networks, which share the same set of nodes (user accounts). A user who is interested in the content published by another account may "follow" it, and thus subscribe to their posts, or "tweets". Over time, each user constitutes their own portfolio of such subscriptions: this defines the first network, the *follower network*. Twitter exposes users to a portion of the tweets published by the accounts they follow, also denoted as "followees". Among this information feed, users may sometimes republish a post that they find particularly relevant, i.e., "retweet" it. This creates the basis for the second network, the *retweet network*, whose links denote the fact that a user retweeted a post that one of their followees previously published (be it an original post or already a retweet).

As such, a follower network denotes potential influence while a retweet network describes some form of actual influence: user retweets reveal successful exposure to content flows enabled by follower links. Furthermore, a follower network may be defined in a static manner. While it evolves when users add (or remove) subscriptions at a given pace, at any time point t a user still has a given and well-defined list of followees and followers. By contrast, a retweet network is fundamentally dynamic: it necessarily stems from the aggregation of observations of retweets over a certain time period and may only be defined by specifying a time range $[t, t']$. More formally, we define:

- the **follower network** F_t at t by adding a directed link $u \rightarrow v$ if u follows v at t , representing potential attention of u to v (as schematically shown in Figure 1a,b left panels). The out-degree k_u of u in that network directly denotes the number of followees of u , while the in-degree k'_v denotes the number of followers of v .
- the **retweet network** $R_{[t,t']}$ over $[t, t']$ by focusing on links $u \rightarrow v$ in F_t , then counting the number of times u retweeted v 's tweets or retweeted a tweet after v published that tweet, over the time period $[t, t']$ —in what follows, this is precisely what we mean by "retweet". We add a weighted directed link $u \rightarrow v$ in $R_{[t,t']}$ with a weight w_{uv} equal to that count (demonstrated in Figure 1b right panels). The out-degree κ_u denotes the number of users whom u retweeted while the in-degree κ'_v denotes the number of users who retweeted v . Distributions of these quantities are shown in Figure 2a. The out-strength s_u denotes the sum of the weights of the out-going links from u , i.e., number of retweets u made of their followees, while the in-strength s'_v denotes the total number of times v has been retweeted by their followers.

Thus, links in $R_{[t,t']}$ form a subset of the links found in F_t . This implicitly relies on the assumption that F_t provides a good approximation of F_τ at another time $\tau \in [t, t']$, which is indeed acceptable if t'

is sufficiently close to t as will be the case in this paper: while the follower network seems to be highly dynamic in the long term, its evolution may be considered to be relatively limited over several weeks, where the proportion of replaced links remains in the vicinity of 10% [31].

Figure 1 illustrates the construction process of both networks, where link directionality denotes some form of attention of the origin to the target. From our data we constructed a directed follower network with 905,112 customers as nodes and 69,156,298 follower relationships as links. The directed retweet network features 428,404 nodes and 937,242 links.

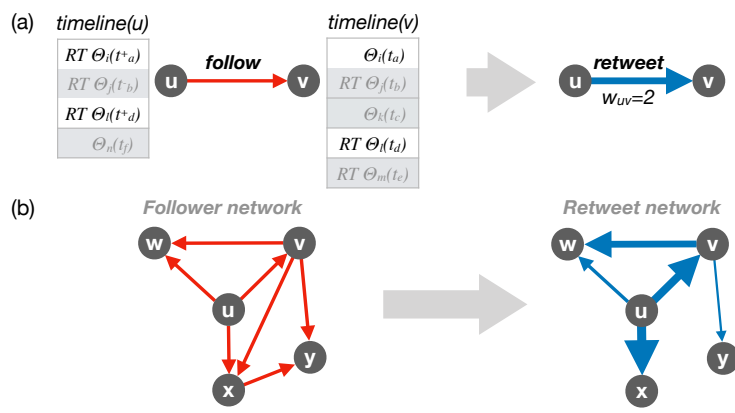


Figure 1. (a) Construction of the retweet network $R_{[t,t']}$. In the timelines of u and v , we distinguish tweets Θ from retweets $RT \Theta$. Given that u follows v , we consider that u retweeted v if u retweeted a tweet or a retweet of v at a posterior time point. In this example, this includes two retweets: Θ_i , which has been posted by v at t_a and retweeted by u at $t^+_a > t_a$, as well as Θ_l , which has been retweeted by v at t_d as well as u at $t^+_d > t_d$. However, this does not include Θ_j , which u retweeted at $t^-_b < t_b$. Other tweets and retweets not common to both timelines are naturally ignored. The weight of the retweet link from u to v is thus $w_{uv} = 2$. (b) The retweet network $R_{[t,t']}$ is therefore a weighted sub-network of F_t : in both networks, the link direction from u to v is aligned with observed attention from u to v .




Figure 2. Distributions of degrees, activity, and similarity measures. (a) Distributions of the κ_u out- and κ'_u in-degrees of nodes in the retweet network. (b) Distributions of the number of retweets (n^{RT}) and number of hashtags (κ^S) per user. (c) Cumulative distribution of the Jaccard similarity between hashtag sets of connected (blue) and randomly selected (orange) pairs of users of the retweet network.

3.3. Attentional Degree

Follower links thus denote a binary notion of potential attention: either one follows, or not. By contrast, retweet links correspond to the magnitude of some form of actual attention. In other words, through retweets we observe how followers diversely allocate part of their attention to their followees. Thus, for a given user u , the distribution of weights of their actual retweets of their followees v_1, \dots, v_{κ_u}

may be heterogeneous to some level: some users may allocate most of their attention to a certain followee, while others may balance their retweets across their whole portfolio of followees. To capture a notion of attention allocation, we use a measure of statistical dispersion, the Herfindahl–Hirschman index (HHI) that we apply on normalized weights of retweet (out-going) links from u :

$$H_u = \sum_i^{\kappa_u} \left(\frac{w_{uv_i}}{s_u} \right)^2 \quad (1)$$

We call *attentional degree* $a_u = 1/H_u$, which varies between 1 (when attention is entirely focused on a single neighbor: $\exists v, w_{uv} = s_u$) and κ_u (when attention is evenly split among all neighbors: $\forall i, w_{uv_i} = 1/s_u$). In economics, for instance, this value is computed on market shares of competing firms in a given market [32]. It is interpreted as the number of “equivalent firms” to assess whether there is sufficient competition. If two firms command each with slightly less than 50% of the market, although many other firms share the rest, this index is going to be close to two, indicating a duopoly despite an apparently high number of firms. In our context, a_u represents by analogy the number of equivalent users whom u devotes a *meaningful* share of their attention to. We contend that this value realistically captures the number of neighbors with whom a node principally communicates in a directed weighted network. Please note that from this value, it would be possible to compute a reduced network by conserving edges with top weights in such a way that the number of neighbors of u would coincide with a_u . It would however require us to additionally specify what to do with edges of equal weight close to the threshold induced by a_u . See also Appendix A for a broader discussion on the meaning of this measure from the viewpoint of network reduction techniques.

We thus consider three layers of attention: first, the follower out-degree as the potential attention, second, the retweet out-degree as the actual attention, third, the attentional degree as the meaningful attention. We are aware that the observation of retweets and the way they are relayed within the follower network remains a partial proxy to describe attention dynamics that does not capture the whole picture of either influence or attention—it misses information, which not only impacted users yet did not lead to observable retweets, but also which circulated through many other possible information channels.

4. Social Attention

4.1. Distribution of Roles

We provide an introductory outlook on the data we analyse and its various basic metrics in Figure 2, in terms of hashtags, retweets, or followers. In particular, the number of both hashtags or retweets per user, and the number of both followers or followees for each user all follow a typical and unsurprising heterogeneous distribution: many have little, few have a lot. We obtained equally unsurprising results when measuring the semantic similarity between pairs of users as the Jaccard index of their hashtag sets. As shown in Figure 2c this similarity appears to be significantly higher for pairs of users connected in the retweet network, compared to random pairs.

In other words, on Twitter, as in many other online interaction contexts, we generally observe homophilic connection patterns while a significant part of the attention is concentrated on a few users. Figure 3 focuses on the retweet network and further illustrates the *distribution of roles* in terms of potential input/output and actual input/output flows (measured as retweets). In particular, panel Figure 3a shows the actual balance of flows as the $P(\kappa'_u/\kappa_u)$ distribution of the ratio between κ'_u the number of times a user is retweeted and the κ_u number of times a user retweets a neighbor in the retweet network. We call this the *retweet balance* and show its distribution as a density plot due to discretization effect. We may similarly define the *follower balance* as the ratio of followers vs. followees in the follower network.




Figure 3. Distributions of roles. (a) The $P(k'_u/k_u)$ distribution of retweet balance across all users shown as a density plot. (b) Configuration of the retweet balance (y-axis) in regard to the follower balance (x-axis), i.e., actual vs. potential attention flows measured as the correlation between the ratio of in- and out-degrees in the retweet and follower networks respectively. Heat-map colors code absolute counts.

The right panel of Figure 3 then provides a broader picture by showing the density of users who exhibit a certain retweet balance (y-axis) with respect to their follower balance (x-axis), i.e., by comparing actual vs. potential attention flows in both directions. This results in a double dichotomy that resembles what had been found earlier by [33], even though they compared the follower with the mention networks. By contrast, we interpret this configuration from an attentional viewpoint, focusing on influence only and differentiating static from dynamic phenomena. In the top right quadrant, we find strong influencers who benefit from an excess of attention, both statically and dynamically; in the bottom left quadrant, what we may call normal users who pay static and dynamic attention to others rather than the other way around. The other two quadrants, top left and bottom right, are comparatively much less populated, they correspond respectively to users with a strong retweet balance yet a weak follower balance (actual flows are way above potential flows, in relative terms: so-called “hidden influentials”), and to users with a weak retweet balance yet a strong follower balance (actual flows are way below potential flows, whom we may denote as “fake influentials”).

These observations contribute to understand Twitter as a decentralized system of users who pay attention to others and forward information in a heterogeneous and somewhat homophilic manner. In this context, we now analyze the traces of these information dynamics to appraise whether cognitive limitations may nonetheless impose constraints on the functioning of this social information system.

4.2. Attention Concentration

To this end, we turn to the measure of meaningful attention, through the analysis of the attentional degree a_u measured in our Twitter data. As shown in Figure 4a, this quantity appears to be heterogeneously distributed, but it takes values over a limited range as compared to the corresponding out-degree distribution of the retweet network. Next, to see how concentrated attention is, we plot the density of a_u vs. k_u in Figure 4b. This tells us to what extent users focus on a certain number of users among their followees. Irrespective of the number of followees, we observe that the attentional degree remains concentrated around a relatively narrow core in the order of about ~ 10 users. For any user, it does not go above a hard threshold of a hundred users, even for users who follow hundreds or thousands of accounts. The observation that the reduced number of meaningful neighbors appears with such an upper limit is in a striking accordance with earlier suggestions. Dunbar’s hypothesis suggests [1,2] that due to cognitive limitations, a person can maintain only about 150 meaningful social relationships at the time (shown as a red solid line in Figure 4b,c). Furthermore, other empirical studies, using different metrics than here, confirmed this suggestion in case of mobile phone communications [4] and in case

of online social platforms, like Twitter [14], which would otherwise allow for more economic ways of interactions than traditional communication means.

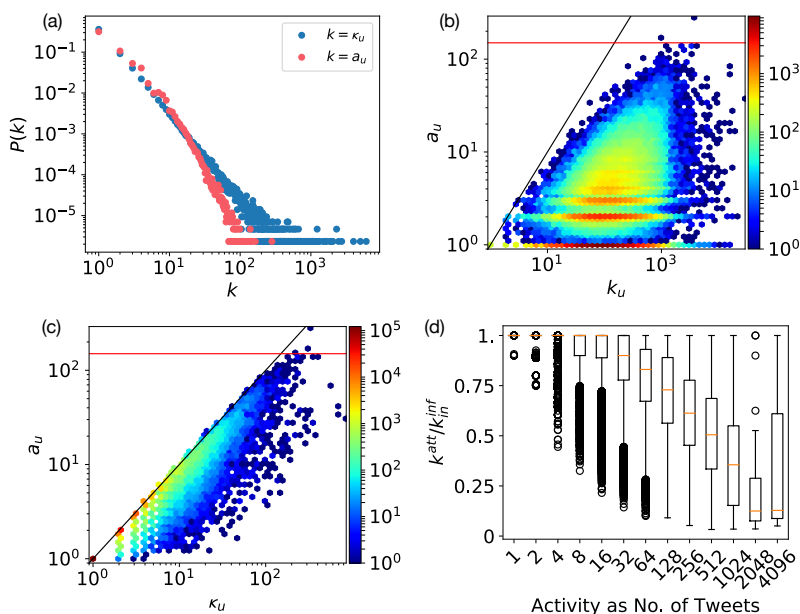


Figure 4. (a) The $P(k)$ distributions of the $k = a_u$ attentional degree and the $k = \kappa_u$ out-degree of the retweet network. (b) Attentional degree vs. out-degree in follower network. (c) Attentional degree vs. out-degree κ_u in the retweet network. (d) Ratio between attentional degree and retweet out-degree κ_u/k_u (y-axis) as a function of activity as the total n number of tweets and retweets (x-axis). All heat-map colors code absolute counts.

We characterize further this core by comparing a_u to κ_u . Deviations from the diagonal indicate to what extent the meaningful number of retweeted users (a_u) is smaller than their total number (κ_u). We plot this Figure 4c, which shows that the two values are somewhat related: density is higher close to diagonal. In other words, there is generally a good correlation between the raw number users who are given any attention to and the equivalent number of such users: once we pay attention to some users among our followees, attention appears to be relatively evenly distributed across them ($a_u \sim \kappa_u$). We propose to describe this configuration as two levels of attention, whereby user first pay most of their attention to a core of their followees while somewhat neglecting a periphery, and then equally pay attention to this core in weighted terms.

Two-level flows of attention. Furthermore, Figure 4c also exhibits some strong deviations, which indicate that some users further restrict their attention to an even smaller super-core of accounts. To understand better who these users are, we examine the ratio between a_u and κ_u and differentiate it with respect to increasing classes of activity, measured as $n = n^{TW} + n^{RT}$, i.e., the total number of tweets and retweets published by a user. Corresponding results in Figure 4d seem to suggest that users with the highest level of activity correspond to those where the deviation regarding the diagonal $y = x$ is strongest, i.e., where there is a higher loss/dissipation between attentional and retweet out-degrees. In other words, beyond a certain level of tweeting activity, only a core of users may receive significant attention, at least in relative terms. The above-described attentional thresholding is thus stronger for more active users, even though they naturally appear to pay attention to a higher raw number of users. This is consistent with similar observations on the allocation of attention on Facebook [17] where the more active users still devote a large portion of their interactions to a small core of about 15 top users.

To summarize, we may describe these dynamics as “two-level flows of attention”: first, users retweet i.e., focus on some followees only, among all potential followees (attentional degree is capped with respect to the number of followees, but generally close to the out-degree in the retweet network); second, higher activity induces a stronger focus in relative terms or, put the other way around, higher activity does not really seem to enable a corresponding widening of attention toward a proportionally more varied set of users. On the whole, this appears to point at the existence of a (converging) social lens which activity does not weaken much.

5. Semantic Attention

5.1. Semantic Attentional Degree

Users not only pay attention to some users, but they may also focus on some issues: their posting activities may be devoted to a variety of topics, which may also be capped. Does informational attention exhibit similar features as interactional attention, is there a link between social attentional constraints and semantic ones? To appraise this, we measure semantic attention in terms of the diversity of hashtags used in a user’s retweets. For a given user u , we consider all their retweets and compute the vector of occurrence of hashtags that they used strictly more than once (i.e., we ignore hapaxes for a given user level): ω_{uh} denotes the number of times a hashtag h has been used by u in their retweets. We may compute the HHI on ω and thus the semantic attentional degree a_u^s as the inverse, which provides an indication of the number of equivalent hashtags or topics addressed by u . Figure 5c exhibits a good correlation between a_u^s and attentional degrees $a_u^{s,all}$ computed on all tweets, not only retweets and, more broadly, all figures of this section were also computed by considering tweets. They all yielded qualitatively similar results, indicating that tweet and retweet behaviors are generally consistent with one another as regards semantic attention.

We compare this with the raw number of hashtags ever addressed by u in their retweets, which we define as the semantic degree κ_u^s by simple analogy with the social degree. Figure 5a features the distribution of a^s vs. κ^s , which is similar to comparing the attentional social degree with the out-degree in the retweet network (even if the results are not as detailed as in the social case, where we can additionally distinguish potential attention from the follower network). Here too, we observe that: first, semantic attention is capped to several hundreds of hashtags in raw terms (κ^s), and to slightly above a hundred topics in equivalent terms (a^s); second, certain users do focus on hashtags in the sense that there is a more or less pronounced deviation between κ^s and a^s . On Figure 5b we show that this dissipation is also strongly dependent on activity, perhaps in a less pronounced manner than in the social case: the semantic attention for the more active users (from an activity of about a hundred tweets) tends to exhibit a magnifying effect that corresponds at most to half the raw number of hashtags. In other words, active users both have broader interests but also display expertise patterns. On the semantic as well as the social sides, these effects exhibit a strong variance. It is likely that on the whole, they indicate the existence of two distinct sub-populations of users: users, at all activity levels, who have a naturally narrow range of interests/interactions, and users, generally among the most active ones, who are broad yet remain quite focused.

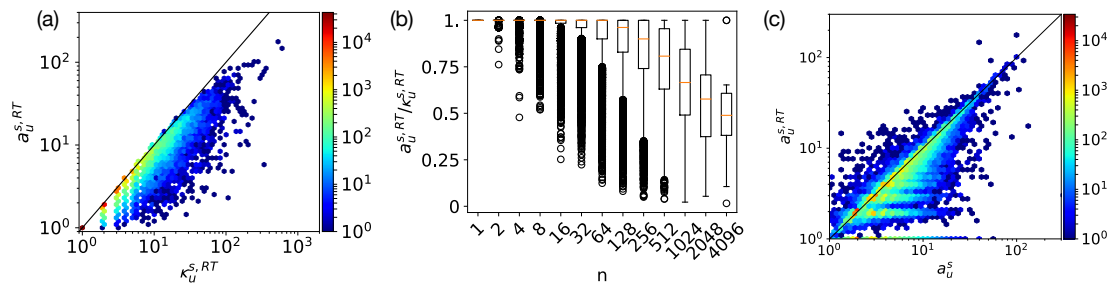


Figure 5. (a): Semantic attentional degree as a function of the hashtag set size. (b): Ratio between semantic attentional degree and number of hashtags a_u^s / κ_u^s (y-axis) as a function of activity as the total n number of tweets and retweets (x-axis). (c) Correlation between semantic attentional degrees a_u^s computed by considering retweeted hashtags, and $a_u^{s,all}$ computed by considering any tweeted or retweeted hashtag of a user u . This correlation appears with a Pearson coefficient $R = 0.889$ (p -value < 0.05). All heat-map colors code absolute counts.

5.2. Socio-Semantic Correlations

Cognitive constraints may thus apply on the interactional and informational sides. We now examine the combination of both and, in particular, aim to verify whether one has a link with the other. Two competing hypotheses may be proposed here. One is that of a joint reinforcement, or at least a positive correlation between both: users who pay attention to more actors also pay attention to more topics. The other hypothesis corresponds rather to a zero-sum game, where attention allocated to one dimension would likely constrain that allocated to the other.

In the former case, i.e., reinforcement, we may hypothesize that there exists an underlying activity variable (here, on Twitter), which would manifest itself in both dimensions: if users are able to cover more interactions, they also cover more topics, and vice-versa. This would be analogous to what has long been observed when comparing online and offline social capitals: while some have suggested that online sociability could deplete the potential for offline sociability, it has been shown that users who are socially more active online are also more active offline [6]. In the latter case, i.e., a zero-sum game, we would observe a negative dependency between social and semantic attention, very much like what is behind Cobb–Douglas consumption graphs in economics where, for a constant level of possible consumption C , two possible goods A and B are consumed according to a function $C = A^\alpha B^{1-\alpha}$ such that consuming A reduces B , generally in a non-linear fashion.

We plot the relationship between a and a^s in Figure 6a–d, distinguishing various levels n of posting activity (i.e., again, in terms of total publications). For all levels, semantic attention appears to be correlated with social attention, with positive R Pearson correlation coefficients and p -values summarized in Table 1. On heatmaps, highest densities are found around the diagonal, while boxplots again confirm a generally positive association between both types of attention. Besides, as said before, there seems to be a non-linear relationship between posting activity and the average value of the social and semantic attentional degrees: centers of mass of degrees on all heatmaps do not move as quickly to the right as the center of mass for the respective retweet number ranges.

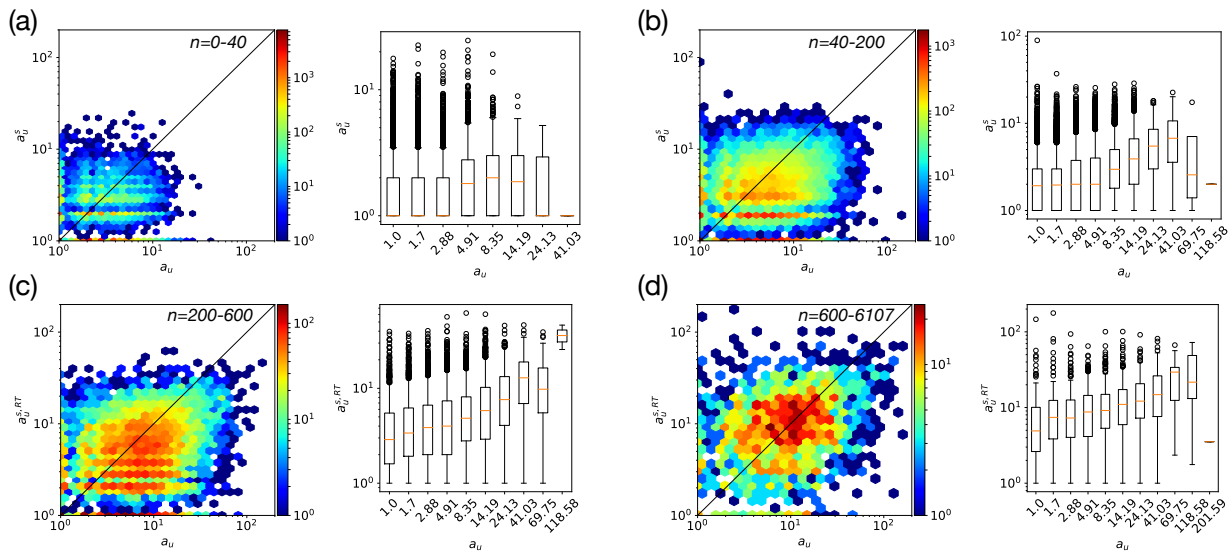


Figure 6. Panels (a–d): semantic attention degree vs. social attention degree, split by the number of retweets. Left panels depict correlation heatmaps between the semantic and social attention degrees of users belonging to a certain activity group. Right panels show the same information as boxplots. Activity n of users are defined as the total number of their tweets and retweets.

Table 1. Pearson correlation coefficient and p-value computed between the social and semantic attentional degrees of individuals in different n retweeting activity groups.

| n | [1, 40[| [40, 200[| [200, 600[| [600, 6107] |
|-----|---------|-----------|------------|-------------|
| R | 0.137 | 0.253 | 0.309 | 0.238 |
| p | <0.05 | <0.05 | <0.05 | <0.05 |

Overall, this seems to generally go in favor of the reinforcement hypothesis, moderated by posting activity in a non-linear fashion. Looking closely at all heatmaps, however, reveals that there is a bright horizontal (resp. vertical) band of high density of hexagons for small values of the vertical (resp. horizontal) axis, i.e., for a given small semantic attentional degree, for instance, there is a horizontal band of bright colored hexagons spanning several orders of magnitude of social attentional degree. To summarize, there seems to be both a strong mass of points loosely around the diagonal, and a strong mass of points along vertical (resp. horizontal) lines for small values on the horizontal (resp. vertical) axis. In other words, some users seem to focus exclusively either on the social side or the semantic users. It is unclear what the status of these users are (especially in terms of them being humans or bots) and this would warrant further research.

6. Limitations

Although we aimed to rely on some of the most general data filtering methods to obtain representative samples of Twitter activities, and on some of the least specific measures to capture attention, our study still has certain limitations. We wish to discuss some of them in this section to make it easier to draw more precise conclusions from our results.

First of all, we measured social attention via retweeting activity of a given neighbor. As discussed earlier, focusing on retweets may provide a partial view of attentional processes. Furthermore, Twitter may offer other mechanisms to quantify more precisely these effects, for example by using *likes*. However,

this type of information was not available to us during the data collection period. Thus, we could not use them for a more precise quantification of attention. Attention signals may also be collected from external platforms: for instance, [34] uses audience data from *bit.ly*, a link tracker website, to study the impact of shared links as a function of the number of followers of the users who posted them. While being a very sound protocol to study conversion rates, it also makes it difficult to match individual user characteristics across both datasets and thereby to discuss user-centric features, which are key from an attentional viewpoint.

Second, Twitter does not only involve human actors but several fake accounts and robots, which may bias our observations. If these non-human actors, or some of them, exhibit unrealistic posting and sharing activities, they may appear as outliers in our measurements—assumedly, they would be unlikely to influence much the overall trends and core observations that we made on a larger population.

Most of the data collected from a Twitter stream come as a sample restricted by pre-defined filters and collection rate limits. While filters are set up by the collector, rate limits induce some ambiguity in the data collection process. Collected data are commonly assumed to represent an unbiased sample of the Twitter stream coming as a fraction of tweets uniformly sampled from the set of tweets meeting the filtering conditions at a given time. While this assumption has been made in almost all Twitter studies, some work [35] addressed and cautioned about the observational bias induced by the unknown sampler algorithm of Twitter. In the case of our dataset, we applied several language and location filters (as explained in Section 3.1) and obtained a relatively high rate of 15–25% tweets via the PowerTrack API as compared to the Open API with only 1% access. Despite this higher rate of data collection, which may considerably reduce the sampling bias we have in our data, we identify this ambiguity as a potential limitation, which is unfortunately present in the vast majority of other Twitter studies as well.

7. Discussion

Twitter may be seen as a decentralized social information processing platform relying on users as input/output devices who are plugged onto their followees' information feeds, part of which they may or may not decide to dispatch to their followers. This decentralization is not devoid of hierarchy and heterogeneity. From this viewpoint, at the collective level, it features a hierarchical yet roughly dichotomized distribution of roles: some users gather a lot of the potential and actual attention, many pay it, while potential and actual attention are generally correlated. Furthermore, at the individual level, we could hypothesize what we may call a "two-level flow of attention" whereby users first focus their actual attention on a core of their potential attention, then redistribute it in a relatively uniform way within that core. This observation was made possible by the use of a simple attentional focus measure, the attentional degree, which consists of a parameter-free approach to compute a number that may be easily be interpreted and compared with raw measures of numbers of neighbors in a network. On the whole, the limitations and focus effects that we find are consistent and, more importantly, extend the broad picture that has been depicted in other platforms in the literature. An interesting question that remains to be addressed would relate to the articulation between the collective and individual levels and, more precisely, to the description of the structural positions of the actors who gather the core of the attention of their neighbors, and the corresponding correlations. For instance, do the peers who are paid the most attention to at a user-centric level also occupy certain positions in the network, do they act as opinion leaders as suggested by Lazarsfeld's two step flow of communication hypothesis [36,37], and are their topological properties correlated among each other? This would shed light on the possible existence of a higher attention likelihood for some kinds of peers who are more likely to pass information and who may be found in specific parts of the network (e.g., in terms of distant communities, cohesive clusters, and so-called hubs and bridges) [38].

Finally, we completed the understanding of the constraints that apply to individual information processing by adopting a joint interactional and informational perspective. In particular, our last figure sheds light on a major question regarding the supply of attention: do semantic and social activities share the same limited supply and thus have a negative impact on one another (convex relationship between both) or do they actually reinforce one another (positive correlation along the diagonal) while being a sublinear function of the activity level? In this respect, we observed a relationship between social and semantic attentional processes whereby they are also generally correlated: we could demonstrate that for most users, both types of attentional resources are related in a positive manner. This hints at the existence of a heterogeneous distribution of attentional resources among users that expresses itself jointly on the semantic and social side, even though there are sometimes marked discrepancies between both types of attention (in terms of divergence regarding the diagonal “social attention=semantic attention”) and that there even exists a special minority of users who are exclusively focused on one side only (i.e., either semantic or social). Here again, a further research direction could consist of appraising the structural positions occupied by these very users—especially by qualifying which users exhibit more social than semantic attention and whether they possess specific topological properties and populate specific parts of the whole system. This, together with the above-mentioned phenomena regarding the uneven distribution of attention, would be likely beneficial to influence studies and contribute to develop finer models that take into account the differential balance of attention among users and towards some selected peers.

Author Contributions: All authors participated in the design of the presented research and in the drafting the publication. A.B. performed the exploratory data analysis. M.K. performed the advanced data analysis. C.R. and M.K. wrote the manuscript.

Funding: This paper has been partially realized in the framework of the “Algodiv” grant (ANR-15-CE38-0001) funded by the ANR (French National Agency of Research). It was also supported by the SoSweet ANR project (ANR-15-CE38-0011), the ACADEMICS grant of the IDEXLYON, project of the Université de Lyon, PIA operated by ANR-16-IDEX-0005, and the “Socsemics” Consolidator grant funded by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 772743).

Acknowledgments: We are grateful for the data collection and preparation to J.-Ph. Magué.

Conflicts of Interest: The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

Appendix A. Backbone Networks and Attentional Degrees

Network backbones [39] provide a filtering method (called the disparity filter) to extract the relevant connection backbone in weighted networks, keeping edges which represent statistically significant deviations for the local assignment of link weights as compared to a null model. In our case we refer to them as a reference system to show that our network reduction method does not evidently capture the same information and thus provides a novel way to identify important ties in a weighted network.

Taking the normalized weight $p_{ij} = w_{ij} / \sum_j w_{ij}$ of links connected to node i we compare their values to a null model where p_{ij} is sampled from a random uniform distribution of the total strength $\sum_j w_{ij}$ of the given node i . This sampling process corresponds to placing $k - 1$ points uniformly over the range of $[0, 1]$ which provides us k sub-interval, where k being the degree of node i . The probability distribution of the value x of each of these intervals [39] is

$$\rho(x)dx = (k - 1)(1 - x)^{k-1}dx. \quad (A1)$$

The disparity filter defines a probability α_{ij} for each link to decide whether it is compatible with the null hypothesis. It is equal to the probability that one of the k intervals defined in the null model has a length larger than p_{ij} , that is

$$\alpha_{ij} = \int_{p_{ij}}^1 \rho(x) dx = (1 - p_{ij})^{k-1}. \tag{A2}$$

In case the link appears with $\alpha_{ij} < \alpha$, the null hypothesis is rejected at the significance level of α and the link is kept in the filtered network. The significance parameter α is the only parameter of the filter and controls the number of links kept in the retained network. Note that in case $k = 1$ for a given node, its single link is automatically preserved.

Here we apply the disparity filter on the directed weighted retweet network to obtain the backbone structure of incoming links, which is corresponding to the reduced attention network of the same structure. This means [39] that while computing p_{ij} we only consider weights of incoming links of node i and in turn $\alpha_{ij} \neq \alpha_{ji}$.

To compare the backbone network and attentional degrees obtained from the retweet network, first we look for the case when they are the most similar in terms of their degrees. To find the best correspondence, we take the attentional degrees and compute the correlation with degrees of several backbone networks obtained by varying the α disparity filter parameter. The highest correlations has been found for $\alpha = 0.575$ as shown in Figure A1a, where we plot the $R(k_{in}^{att}, k_{in}^{bb})$ Pearson correlation coefficient between node degrees in the two structures as the function of α . Please note that all observed correlations are significant with $p < 0.05$. Comparing the degree distribution of the retweet network to the similar filtered graphs (see Figure A1b) we find that although they are very similar, the attention filter provides a network with the most reduced degree heterogeneities. To directly observe degree correlations, we show the degrees of nodes in the two selected filtered structures as a heat-map of a scatter plot in Figure A1c. There, despite the strong correlation between the various degree values ($R = 0.955, p < 0.05$), we find strong fluctuations as well, indicating that the two filtering process do not identify the same set of links to be important, which underlies the relevance of our method. Please note that while our method is parameterless, disparity filter has a parameter, what we tuned to obtain the most similar structure. Without this tuning our method may provide even more different filtered set of links, which may hold different roles in the structure. Inversely, attentional degrees can be used to identify the optimal α disparity filter parameter without looking at more complicated network characteristics, as it was suggested originally [39].

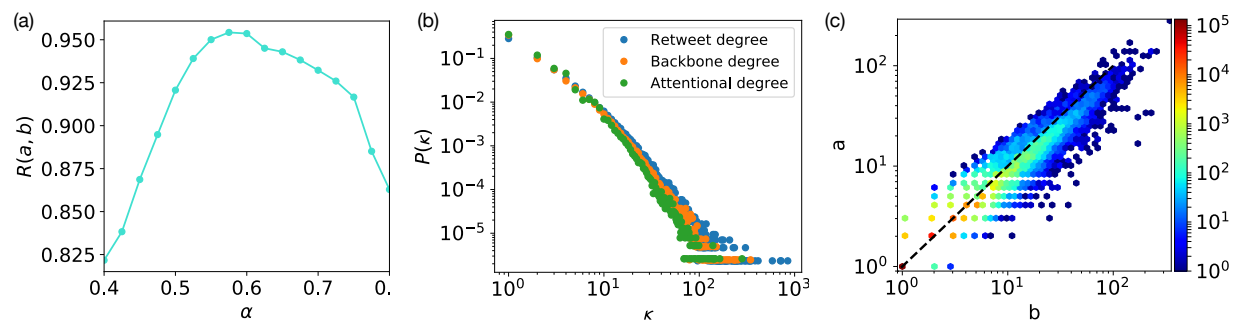


Figure A1. Distributions and correlations of degrees measured in the backbone networks or as attentional degrees. (a) $R(a, b)$ Pearson correlation values between the a attentional and b backbone degrees as the function of the α disparity filter parameter. (b) Distributions of retweet, backbone, and attentional degrees of users. (c) Correlations between the a attentional and b backbone degrees computes with the optimal disparity filter value $\alpha = 0.575$.

References

1. Hill, R.A.; Dunbar, R.I. Social network size in humans. *Hum. Nat.* **2003**, *14*, 53–72. [[CrossRef](#)]
2. Roberts, S.G.; Dunbar, R.I.; Pollet, T.V.; Kuppens, T. Exploring variation in active network size: Constraints and ego characteristics. *Soc. Netw.* **2009**, *31*, 138–146. [[CrossRef](#)]
3. Burke, M.; Kraut, R.; Marlow, C. Social capital on Facebook: Differentiating uses and users. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Vancouver, BC, Canada, 7–12 May 2011; pp. 571–580.
4. Miritello, G.; Lara, R.; Cebrian, M.; Moro, E. Limited communication capacity unveils strategies for human interaction. *Sci. Rep.* **2013**, *3*, 1950. [[CrossRef](#)]
5. Weng, L.; Flammini, A.; Vespignani, A.; Menczer, F. Competition among memes in a world with limited attention. *Sci. Rep.* **2012**, *2*, 335. [[CrossRef](#)]
6. Wellman, B. Computer Networks as Social Networks. *Science* **2001**, *293*, 2031–2034. [[CrossRef](#)]
7. Wellman, B.; Haase, A.Q.; Witte, J.; Hampton, K. Does the Internet Increase, Decrease, or Supplement Social Capital?: Social Networks, Participation, and Community Commitment. *Am. Behav. Sci.* **2001**, *45*, 436–455. [[CrossRef](#)]
8. Orbach, M.; Demko, M.; Doyle, J.; Waber, B.N.; Pentland, A.S. Sensing Informal Networks in Organizations. *Am. Behav. Sci.* **2015**, *59*, 508–524. [[CrossRef](#)]
9. Nguyen, G.D.; Lethiais, V. The Impact of Social Networks on Sociability: The Case of Facebook. *Réseaux* **2016**, *195*, 165–195. [[CrossRef](#)]
10. Newman, M.E.J. Scientific collaboration networks. I. Network construction and fundamental results. *Phys. Rev.* **2001**, *64*, 016131. [[CrossRef](#)]
11. Ebel, H.; Mielsch, L.I.; Bornholdt, S. Scale-free topology of e-mail networks. *Phys. Rev. E* **2002**, *66*, 035103. [[CrossRef](#)]
12. Kumar, R.; Novak, J.; Tomkins, A. Structure and evolution of online social networks. In *Link Mining: Models, Algorithms, and Applications*; Springer: New York, NY, USA, 2010; pp. 337–357.
13. Mislove, A.; Marcon, M.; Gummadi, K.P.; Druschel, P.; Bhattacharjee, B. Measurement and analysis of online social networks. In Proceedings of the 7th ACM SIGCOMM conference on Internet measurement, San Diego, CA, USA, 24–26 October 2007; pp. 29–42.
14. Gonçalves, B.; Perra, N.; Vespignani, A. Modeling Users' Activity on Twitter Networks: Validation of Dunbar's Number. *PLoS ONE* **2011**, *6*, e22656. [[CrossRef](#)]
15. Lehmann, J.; Gonçalves, B.; Ramasco, J.; Cattuto, C. Dynamical Classes of Collective Attention in Twitter. In Proceedings of the 21st International Conference on World Wide Web, Lyon, France, 16–20 April 2012; pp. 251–260.
16. Sanlı, C.; Lambiotte, R. Local variation of hashtag spike trains and popularity in twitter. *PLoS ONE* **2015**, *10*, e0131704. [[CrossRef](#)]
17. Backstrom, L.; Bakshy, E.; Kleinberg, J.M.; Lento, T.M.; Rosenn, I. Center of attention: How facebook users allocate attention across friends. In Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media, Barcelona, Spain, 17–21 July 2011.
18. Saramäki, J.; Leicht, E.A.; López, E.; Roberts, S.G.; Reed-Tsochas, F.; Dunbar, R.I. Persistence of social signatures in human communication. *Proc. Natl. Acad. Sci. USA* **2014**, *111*, 942–947. [[CrossRef](#)]
19. Lin, Y.R.; Keegan, B.; Margolin, D.; Lazer, D. Rising tides or rising stars?: Dynamics of shared attention on Twitter during media events. *PLoS ONE* **2014**, *9*, e94093. [[CrossRef](#)]
20. Golder, S.A.; Huberman, B.A. Usage patterns of collaborative tagging systems. *J. Inf. Sci.* **2006**, *32*, 198–208. [[CrossRef](#)]
21. Wu, F.; Huberman, B.A. Novelty and collective attention. *Proc. Natl. Acad. Sci. USA* **2007**, *104*, 17599–17601. [[CrossRef](#)]
22. Weng, L.; Menczer, F.; Ahn, Y.Y. Virality Prediction and Community Structure in Social Networks. *Sci. Rep.* **2013**, *3*, 2522. [[CrossRef](#)]

23. Romero, D.M.; Meeder, B.; Kleinberg, J. 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, Hyderabad, India, 28 March–1 April 2011.
24. Hodas, N.O.; Lerman, K. How visibility and divided attention constrain social contagion. In Proceedings of the 2012 International Conference on Privacy, Security, Risk and Trust (PASSAT), and 2012 International Conference on Social Computing (SocialCom), Amsterdam, The Netherlands, 3–5 September 2012; pp. 249–257.
25. Rodriguez, M.G.; Gummadi, K.; Schoelkopf, B. Quantifying information overload in social media and its impact on social contagions. *arXiv* **2014**, arXiv:1403.6838.
26. Unicomb, S.; Iñiguez, G.; Karsai, M. Threshold driven contagion on weighted networks. *Sci. Rep.* **2018**, *8*, 3094. [[CrossRef](#)]
27. Suh, B.; Hong, L.; Pirolli, P.; Chi, E.H. Want to be retweeted? Large-Scale analytics on factors impacting retweet in twitter network. In Proceedings of the 2010 IEEE Second International Conference on Social Computing, Minneapolis, MN, USA, 20–22 August 2010; pp. 177–184.
28. Yang, Z.; Guo, J.; Cai, K.; Tang, J.; Li, J.; Zhang, L.; Su, Z. Understanding retweeting behaviors in social networks. In Proceedings of the 19th ACM International Conference on Information and Knowledge Management, Toronto, ON, Canada, 26–30 October 2010; pp. 1633–1636.
29. Zaman, T.R.; Herbrich, R.; Van Gael, J.; Stern, D. Predicting Information Spreading in Twitter. In Proceedings of the Computational Social Science and the Wisdom of Crowds Workshop (Colocated with NIPS 2010), Whistler, BC, Canada, 10 December 2010.
30. boyd, D.; Golder, S.; Lotan, G. Tweet, tweet, retweet: Conversational aspects of retweeting on twitter. In Proceedings of the 2010 43rd Hawaii International Conference on System Sciences, Honolulu, HI, USA, 5–8 January 2010; pp. 1–10.
31. Myers, S.A.; Leskovec, J. The bursty dynamics of the twitter information network. In Proceedings of the 23rd International Conference on World Wide Web, Seoul, Korea, 7–11 April 2014; pp. 913–924.
32. Rhoades, S.A. The herfindahl-hirschman index. *Fed. Res. Bull.* **1993**, *79*, 188.
33. González-Bailón, S.; Borge-Holthoefer, J.; Moreno, Y. Broadcasters and hidden influentials in online protest diffusion. *Am. Behav. Sci.* **2013**, *57*, 943–965. [[CrossRef](#)]
34. Gabelkov, M.; Ramachandran, A.; Chaintreau, A.; Legout, A. Social clicks: What and who gets read on Twitter? In Proceedings of the 2016 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Science, Antibes Juan-les-Pins, France, 14–18 June 2016; Volume 44, pp. 179–192.
35. Morstatter, F.; Pfeffer, J.; Liu, H. When is it biased?: Assessing the representativeness of twitter’s streaming API. In Proceedings of the 23rd International Conference on World Wide Web, Seoul, Korea, 7–11 April 2014; pp. 555–556.
36. Berelson, B.; Gaudet, H.; Lazarsfeld, P.F. *The People’s Choice: How the Voter Makes up his Mind in a Presidential Campaign*; Columbia University Press: New York, NY, USA, 1968.
37. Katz, E.; Lazarsfeld, P.F.; Roper, E. *Personal Influence: The Part Played by People in the Flow of Mass Communications*; Routledge: Abingdon, UK, 2017.
38. Weng, L.; Karsai, M.; Perra, N.; Menczer, F.; Flammini, A. Attention on weak ties in social and communication networks. In *Complex Spreading Phenomena in Social Systems*; Springer: Cham, Switzerland, 2018; pp. 213–228.
39. Serrano, M.Á.; Boguñá, M.; Vespignani, A. Extracting the multiscale backbone of complex weighted networks. *Proc. Natl. Acad. Sci. USA* **2009**, *106*, 6483–6488. [[CrossRef](#)]

