

Distinguishing between Topical and Non-topical Information Diffusion Mechanisms in Social Media

Przemyslaw A. Grabowicz
MPI-SWS, Germany

Niloy Ganguly
IIT Kharagpur, India

Krishna P. Gummadi
MPI-SWS, Germany

Abstract

A number of recent studies of information diffusion in social media, both empirical and theoretical, have been inspired by viral propagation models derived from epidemiology. These studies model the propagation of memes, i.e., pieces of information, between users in a social network similarly to the way diseases spread in human society. Importantly, one would expect a meme to spread in a social network amongst the people who are interested in the topic of that meme. Yet, the importance of topicality for information diffusion has been less explored in the literature.

Here, we study empirical data about two different types of *memes* (hashtags and URLs) spreading through the Twitter's online social network. For every meme, we infer its topics and for every user, we infer her topical interests. To analyze the impact of such topics on the propagation of memes, we introduce a novel theoretical framework of information diffusion. Our analysis identifies two distinct mechanisms, namely topical and non-topical, of information diffusion. The non-topical information diffusion resembles disease spreading as in *simple contagion*. In contrast, the topical information diffusion happens between users who are topically aligned with the information and has characteristics of *complex contagion*. Non-topical memes spread broadly among all users and end up being relatively popular. Topical memes spread narrowly among users who have interests topically aligned with them and are diffused more readily after multiple exposures. Our results show that the topicality of memes and users' interests are essential for understanding and predicting information diffusion.

1 Introduction

The process of information diffusion is crucial for viral marketing (Leskovec, Adamic, and Huberman 2007), spreading of political movements (Strang and Soule 1998) and innovations (Rogers 1962), and opinion formation (Holley and Liggett 1975). Numerous studies have shown that information and behaviors propagate virally over online social networks via *social exposures*, defined as the number of friends or acquaintances of a user that have exposed the meme to her (Backstrom et al. 2006; Cosley, Huttenlocher, and Kleinberg 2010; Crandall et al. 2008; Gomez-Rodriguez, Gummadi, and Schölkopf 2014; Romero, Meeder, and Kleinberg

2011). As such, the phenomenon of information diffusion in social networks is predominantly modeled by building parallels with the viral spread of diseases (Kempe, Kleinberg, and Tardos 2003), or by using graph-based methods (Granovetter 1978; Kempe, Kleinberg, and Tardos 2003). However, information diffusion involves intelligent individuals exhibiting complex behaviors, hence one can expect that it includes mechanisms beyond viral spreading via social exposures. An important mechanism that may drive such spread is the topical alignment between the interests of individuals and the pieces of information, which in the language of epidemiology is expressed as the genetic susceptibility to a certain disease (Burgner, Jamieson, and Blackwell 2006).

However, it is difficult to empirically observe this phenomenon for information diffusion. First, large-scale datasets of information that are not influenced by recommender systems and contain enough information for topic modeling are scarce. Second, topic models have matured and have become viable for large datasets only in recent years. We overcome these issues thanks to the increasing availability of data characterizing users and content in online social networks (Cha et al. 2010) and the development of reliable topic modeling methods (Blei, Ng, and Jordan 2003; McCallum 2002). Third, information diffusion within a given system may be caused by the exposures external to that system, which we are not aware of. Thus, it is challenging to isolate the effect of the external exposures from the effect of the internal social exposures on the information diffusion (Bakshy et al. 2012).

Here, to separate internal exposures from external exposures and to analyze the effect of topics on information diffusion, we introduce a theoretical framework that models topics of memes and users as latent variables and the probability of diffusion as a function of these variables. Diffusion processes are usually described statistically with the so-called *adoption probability* (Backstrom et al. 2006; Cosley, Huttenlocher, and Kleinberg 2010; Crandall et al. 2008; Gomez-Rodriguez, Gummadi, and Schölkopf 2014; Romero, Meeder, and Kleinberg 2011), defined as the probability of adopting a meme by a person who has been exposed to this meme by exactly κ social exposures. Our theoretical framework allows us to extend the definition of adoption probability by accounting for the topical alignment between a meme and interests of each user and to estimate the infor-

mation diffusion from both internal and external sources.

Furthermore, we hypothesize that there exist two different mechanisms of information propagation, namely, a topical and non-topical information diffusion. Under this hypothesis, topical information would start to spread from the enthusiasts or experts of that topical domain. Then, it would diffuse narrowly to other members of this domain. Non-topical information would not be limited by its topical domain. If such hypothesis is correct, then one can expect that topical and non-topical information diffusion have different properties. Here, we aim to identify characteristics specific to each of these two mechanisms. Finally, having the two mechanisms identified and characterized, it is crucial to relate them to existing theories of information diffusion. The sociological theory of complex contagion suggests that multiple social exposures are important for the diffusion of novel and controversial information (Centola and Macy 2007), while the theory of simple contagion posits viral spreading after the initial exposure. How these theories relate to topical and non-topical information diffusion? These are some of the questions that we address in this study.

1.1 Present work

Note that, to understand the importance of topicality for information diffusion, we require (i) a closed social system that does not implement any personalized recommendation algorithm¹, (ii) a comprehensive and independent annotation of topicality of users and diffused pieces of information, and (iii) a theoretical framework that is able to estimate the effect of internal and external social exposures on information diffusion and to quantify topicality of memes and users.

To this end, we analyze a dataset from Twitter containing all public tweets that were messaged during the observation period of three months in mid-2009 and a full snapshot of the follower graph from the end of that period. During that time, there were no personalized recommendations in Twitter and no other mechanisms than follower links to subscribe to other users. Hence, we are able to measure information diffusion in a *natural setting*, namely in an online social network that is hardly affected by any recommender system. We consider two distinct types of information that diffuse over the follower graph, namely hashtags and URLs, which for simplicity we jointly call *memes*. We say that a user *diffuses* or *adopts* a meme if she writes it in one of her tweets.

Next, we split the dataset chronologically into two parts (see Figure 1). We process the prior part to infer the topics of memes and users’ interests, whereas we use the latter part to perform an independent analysis of information diffusion. More specifically, using nouns extracted from past tweets we infer with an unsupervised method, namely with Latent Dirichlet Allocation, i.e., LDA (Blei, Ng, and Jordan 2003), the topical distributions of active users and popular memes (Section 3). As such, in our theoretical framework, topics are mixtures of words that frequently appear together. Each user and meme is represented by a topical distribu-

tion over a constant number of topics. To evaluate how topically narrow are memes and users’ interests, we compute the entropy of their corresponding topical distributions. Thanks to this procedure, we distinguish between *topical* and *non-topical* memes and users. We say that a meme is topical if it is strongly associated with a topic, whereas a non-topical meme is weakly and equally associated with all topics. Importantly, we define the topical alignment between a meme and interests of a user as a cosine similarity of their respective topical distributions. We exploit this topical alignment to compute the probability of adoptions triggered either via sources *internal* or *external* to Twitter (Section 4). We find that topical memes are likely to be topically aligned with the interests of people who adopt them, either via internal or external exposures. We leverage out theoretical framework further, finding that:

- A considerable fraction of memes are topical. (Section 3)
- We show that adoptions of topical memes are easier to predict than the adoptions of non-topical memes. (Section 5.1)
- For the topical information diffusion, the probability of adoption steadily increases with the topical alignment between user’s interests and a meme, whereas for non-topical information diffusion the probability stays flat or increases considerably less. (Section 5.1)
- Topical information diffusion profits more than non-topical diffusion from multiple exposures to a meme; thus, it resembles complex contagion. We capture it by introducing a novel metric of spreading persistence. (Section 5.2)
- A large fraction of adoptions happens via seed adopters exposed to external sources. The seed adopters of topical memes are relatively more aligned with the meme than other adopters and the seed adopters of non-topical memes. (Section 5.3)
- The diffusion of URLs is about twice more intense via external rather than internal sources, whereas the diffusion of hashtags has similar intensity both for external and internal exposures. (Section 5.3)

1.2 Related work

Some of our results can be interpreted with the sociological theory of *complex contagion* (Centola 2010; Centola and Macy 2007). This theory posits that “high-risk”, “novel”, “unproven”, or “controversial” information spreads via “wide bridges”, namely via multiple social exposures or strong ties. As we show later, topical memes convey such information. Romero et al. report that successive social exposures to political hashtags tend to increase adoption probability more than multiple social exposures to idioms (Romero, Meeder, and Kleinberg 2011). The authors explain these diffusion classes with complex contagion mechanism. Our work generalizes their analysis and methodology, in the sense that we do not focus on a set of specific topics (e.g., politics and idioms), but instead we study a quantifiable and robust topical space, which represents such specific topics, yielding consistent results. Our

¹At the time of data collection, all follower links were created by users themselves without the influence of any personalized recommendation algorithm.

findings suggest that topical information diffusion is related to complex contagion.

It is intuitive that human behavior is largely related to multi-dimensional personal characteristics of people, such as topical interests or private traits (Kosinski, Stillwell, and Graepel 2013). For instance, some of the recently proposed models suggest that latent factors impact information diffusion (Barbieri, Bonchi, and Manco 2012; Du et al. 2013) and measure topical social influence (Weng et al. 2010; Bi et al. 2014). Moreover, recommender systems exploiting the premise of latent factors predicting content consumption are widely deployed (Ricci, Rokach, and Shapira 2011). However, empirical studies of information diffusion in social media tend to focus either on predictions of the size of social cascades (Ma, Sun, and Cong 2013; Weng and Menczer 2014), analysis of the structure of the cascades (Goel, Watts, and Goldstein 2012), or detecting influential spreaders (Cha et al. 2010), which are macroscopic and mesoscopic properties of information diffusion. A few empirical works study the information diffusion at the microscopic level. These studies explore, for instance, the classes of information diffusion for specific topics (Romero, Meeder, and Kleinberg 2011), the impact of activity similarity on the probability of interaction (Crandall et al. 2008), and the predictiveness of limited content-based features for future adoptions (Yang et al. 2012; Yuan et al. 2016). Thus, to our knowledge, our work is the first large-scale empirical study systematically analyzing the relation between topicality and mechanisms of information diffusion and the extent to which topical alignments affect information diffusion.

2 Dataset

We analyze a dataset from Twitter that is nearly complete and contains all public tweets produced by users until September 2009 and a snapshot of the social graph crawled in September 2009 (Cha et al. 2010). In Twitter, a user can post short messages, called tweets, that are delivered to all her *followers*, namely all the users who follow her in the follower graph. (The users who are followed by the user are her *followees*.) Such messages can contain memes, so the memes can diffuse over the follower graph. Importantly, at the time period when our dataset was gathered, there were no other ways than the follower links to directly receive content from other users. We study the diffusion of two types of memes: hashtags and URLs. We say that a user becomes an *adopter* of a meme and starts to spread it once she posts it

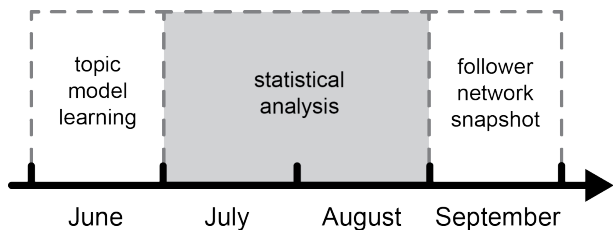


Figure 1: Timeline of the dataset illustrating the design of the study. The gray area marks the period of analysis.

| | |
|------------------|---|
| Iran election | mousavi, basij, khamenei, ppl, protesters, rally |
| British election | bbc, london, bnp, news, brown, bit |
| Well-being | success, happiness, einstein, albert, john, failure |
| Viral marketing | marketing, business, twitter, media, blog, site |
| Smart devices | iphone, apple, app, mac, macbook, wwdc |

Table 1: An example of five topics detected by LDA and words associated with each of them. The words are ordered by their relevance to the topic. (Words “albert” and “einstein” are associated with one of the topics because his sayings are quoted several times within that topic.)

at least in one tweet. Naturally, at the moment of posting the meme, the adopter becomes an *exposer* of the meme to all her followers.

We analyze adoptions that happened during the period of three months, from June to August 2009. During this time users posted nearly a billion of tweets. We consider *only* the memes that emerged during that period, specifically during June. In other words, we do not consider memes that were born before that period. This constitutes a difference from the prior works (Gomez-Rodriguez, Gummadi, and Schölkopf 2014), which tracked also the memes that emerged before the observation period, hence underestimating the number of social exposures.² We focus on memes that were used predominantly in English tweets. Specifically, we analyze the memes for which at least 90% of tweets were in English. To detect the language of tweets we use a state-of-the-art language detection tool for Twitter data (Nakatani 2012). Furthermore, we report results only for the memes that have been adopted by at least 100 unique users.³ Overall, we analyze 704 hashtags adopted 480k times by a total of 330k unique users and 3745 URLs adopted 1.2m times by a total of 610k unique users.

We divide the three-months long dataset into two parts. The first part consists of all tweets from June, while the second part consists of all tweets from July and August (Figure 1). We use the first part of the dataset solely to infer topical distributions of users’ interests and memes emerging during that month. This partition allows us to keep the topic modeling independent from the statistical analysis, which is performed solely on the second part of the dataset.

3 Topicality of memes and users

Here, we describe how we obtain the topical distributions of memes and users’ interests. Then, we introduce the notion of topicality to distinguish between topical and non-topical memes and user’s interests based on their respective topical distributions. Finally, we introduce the topical alignment between a meme and user’s interests.

²The user could have been exposed to the meme before the observation period. If this happens and one does not have knowledge about it, then the real social exposure is higher than the one estimated based solely on the observation period. Our dataset covers all previous months apart from the three-months long period of observation. Thus, we can accurately estimate the real number of exposures.

³We obtain qualitatively the same results for lower thresholds.

The most topical hashtags

| Hashtag | Topic | Examples of tweets hand-picked from 30 chronological tweets from the beginning of June 2009 |
|-----------------|------------------|--|
| #ashraf | Iran election | 1. canadian friends of a democratic iran urge us president to ensure ashraf residents treated according international law #iran #ashraf #rajavi 2. suppressive forces heavily present in streets of iran capital to prevent protests #iran #iranelection #tehran #neda #rajavi #ashraf 3. how u can help iranians http://tinyurl.com/m8fun4 #iranelection #tehran #azadi #ashraf #rajavi #neda #basij #iran #gr88 rt rt |
| #electionstudio | British election | 1. bnp is a left wing party, so were the german national socialists. #electionstudio 2. #electionstudio this bnp guy is a lunatic frankly, indigenous population? if it weren't for immigration britain would not exist at all. 3. #electionstudio i got very excited last night when purnell resigned but now can't see a leadership challenge coming from anywhere |
| #vene2ia | Well-being | 1. without freedom of choice there is no creativity kirk, star trek #vene2ia 2. it wasn't until quite late in life that i discovered how easy it is to say "i don't know!" somerset maugham #vene2ia (via @vene2ia) 3. rt @vene2ia: children are the living messages we send to a time we will not see john w. whitehead #vene2ia |
| #smum | Viral marketing | 1. learn how to blog your way to success at social media university milwaukee http://ow.ly/g83u (last day for early bird pricing) #smum 2. lrt @philgerb: rt @dnewman:great article on benefits of audience tweeting during your presentation http://bit.ly/rdgsl good 2 know 4 #smum! 3. l@backseatgaffer glad you found it. you can sign up even if you're not going to #smum. details: http://bit.ly/3qyhot reg: http://bit.ly/zvpeh |
| #mytouch | Smart devices | 1. mytouch3g ordering starts tomorrow, very excited #tmobile #android #mytouch3g 2. im thinking of getting the tmobile #mytouch phone, i want to know, what do you all think of the phone? should i get it? 3. if you're using a #bold, #curve, #g1, #mytouch or an #iphone you should download @loopt an amazing lbs networking app |

The least topical hashtags

| Hashtag | Topic | Examples of tweets hand-picked from 30 chronological tweets from the beginning of June 2009 |
|---------------|-------|--|
| #uknowyouugly | - | 1. #uknowyouugly when you're at the zoo and the monkeys run from you 2. #uknowyouugly when u gotta get fresh to go to the grocery store 3. #uknowyouugly when they ban you from the mac store |
| #unala2009 | - | 1. getting ready for my first unconference... the first one might as well be a big one! very excited. #ala2009 #unala2009 2. got new business cards designed by my son. unconference tag #unala2009 per conference wiki. ala facebook meetup: http://tinyurl.com/loayhl 3. just for clarification, friday's unconference is at the chicago hilton, not the palmer house hilton. hth #ala2009 #unala2009 |
| #skiphop | - | 1. i just entered to #win a #skiphop bag with a #beautyfix kit in the #99daysofsummer #contest at #modernmom! http://bit.ly/wetvi 2. modernmom i just entered to #win a #skiphop bag with a #beautyfix kit in the #99daysofsummer #contest at #modernmom! http://bit.ly/p29uz 3. i just entered to #win a #skiphop bag with a #beautyfix kit in the #99daysofsummer #contest at #modernmom! http://bit.ly/7ap8x #trackle |
| #sexisbetter | - | 1. #sexisbetter wen your drunk 2. #sexisbetter with a scorpio! (guess what sign i am) :-) 3. #sexisbetter when its anywhere but a bed |
| #besthead | - | 1. #besthead when dey don't expect it baq.. lol 2. #besthead is when u black out and just pass out now dats da best 3. #besthead was in the morning right after she brushed her teeth,she went to work!! u heard me (master p voice) |

Table 2: The top five most and least topical hashtags. For each hashtag, we show three exemplary tweets and a topic corresponding to this hashtag. The least topical hashtags are not associated with a specific topic. In this list, we omit all hashtags that co-appear with the listed hashtags.

3.1 Topical distributions

We define topical distributions for each meme and user based on all tweets posted in June. First, we filter out the tweets that are not written in English. Second, we extract from remaining tweets all nouns and proper nouns using a state-of-the-art POS-tagger for Twitter data (Owoputi et al. 2013). We choose nouns and proper nouns, because they encapsulate the semantics of tweets. We represent each meme as a bag of unique nouns that appeared in tweets containing this meme and each user as a bag of unique nouns that were tweeted by this user. Next, we model the topics of users and memes with LDA performed for 100 topics and 1000 iterations (McCallum 2002).⁴ Hence, the topics of user's interests and a meme are represented by distributions over 100 topics. Examples of topics inferred with this method are shown in Table 1. Note that each such topic gathers related terms. We discuss these topics further in the following paragraphs.

⁴We also tested lower and higher numbers of topics, finding similar results both qualitatively and quantitatively.

3.2 Topicality

To measure how much topical a meme or a user's interests are, we compute the entropy H of their respective topical distributions. The higher the entropy, the less topical the meme or the user are, and vice versa. It follows that the most topical would be a meme that is associated only with a single topic, whereas the least topical would be a meme that is equally associated with all topics, i.e., has a uniform distribution over all topics. Most of memes is between these two extreme examples. The distributions of entropy resemble a normal distribution (Figure 2). We refer to the 25% most topical memes and users as *topical*, whereas to the 25% least topical as *non-topical*. We establish this distinction for the sake of clarity, however, other thresholds (e.g., at 50%) give qualitatively the same results. In the remainder, we present results both for topical and non-topical memes and users.

We list the most and the least topical memes in Table 2. For instance, hashtags #ashraf and #electionstudio correspond to political elections in Iran and Great Britain, #vene2ia represents a motivational website, #smum is about a newly-established viral marketing school, and #mytouch corresponds to a newly released smartphone. Note these

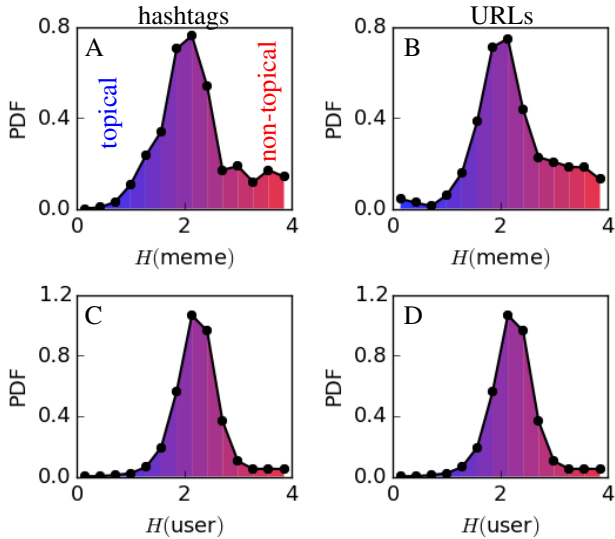


Figure 2: The distribution of entropy of topical distributions of (A-B) memes and (C-D) user’s interests. The memes or Users’ interests with the lowest entropy are the most topical, while the memes with the highest entropy are the least topical.

hashtags correspond closely to the categories of information that is predicted to spread via complex contagion, namely novel, unproven, high-risk, and controversial information.

On the contrary, the least topical hashtags are mostly Twitter idioms and neologism that can be understood directly from their names, i.e., #uknowyouugly, #sexisbetter, #besthead, or are related to a popular lottery, i.e., #skiphop. These idioms have a generic entertainment value. The hashtag #unala2009 corresponds to an un-conference, which is a conference without agenda.

Strikingly, the distinction between topical, e.g., political, hashtags versus idioms corresponds very closely to the categories of hashtags exhibiting complex and simple contagions described in (Romero, Meeder, and Kleinberg 2011) (see Section 1.2). This indicates that the topical information diffusion is related to complex contagion.

Finally, we note that topical memes are $(42 \pm 2)\%$ more likely to be adopted by a topical user than by a non-topical user. This is in a sense the tip of the iceberg, because here we do not take into account topical alignments between memes and users, which increase the probability of adoption much further.

3.3 The topical alignment between a meme and a user

Once we have the topical distributions of memes and users, we measure the topical alignment between a user and a meme, denoted as $S(\text{user}, \text{meme})$, by computing the cosine similarity of the corresponding topical distributions. The topical alignment assumes values from 0 to 1, where 1 corresponds to the maximum possible alignment. The alignment can be 1 both for topical and non-topical memes and users. Even if a meme is perfectly non-topical (has a uniform dis-

tribution over topics), its topical alignment with a user that has interests distributed uniformly over topics will be 1.

4 The theory of adoption probability

In this section, we formally introduce adoption probability and extend its definition to cover topical alignment. To this end, we define the fundamental random events corresponding to the adoptions and exposures of memes. Then, we exploit these events to define the adoption probability as a function of social exposure and topical alignment. Also, we show that the two channels of social influence, namely influence internal and external to Twitter, can be separated within our model of information diffusion. We close this section by introducing a new metric of spreading persistence that quantifies the effect of multiple exposures on the probability of adoption. The methods and definitions that we introduce here can be used for studying information diffusion in various systems. In the remainder of this study, we apply this theoretical framework to our Twitter dataset.

4.1 Extending the definition of adoption probability

One can reason about information diffusion by defining *exposure event* as a tuple (user, meme, exposers, adoption state) that represents a user being exposed by a given set of exposers to a meme that she has not adopted yet. Naturally, the set of exposers may be empty, if the user has not been exposed to the given meme yet by anybody. The set of exposers can only grow in time, as more people adopt the given meme. The adoption state tells us if the meme will be adopted *before* the appearance of a new exposers. If it gets adopted, then we call such exposure event as an *adoption event*, to distinguish it from all other exposure events. Note that each exposure event can be characterized by the topical alignment between a meme and a user. Some users in our dataset do not have topical distributions defined. Thus, we filter out all exposure events for such users.

Next we focus on exposure events with κ exposers. Namely, we say that a user is κ -exposed to a meme at a certain time, if κ distinct users that she follows adopted the meme prior to this time and she has not adopted the meme yet. The probability of adoption is defined (Cosley, Huttenlocher, and Kleinberg 2010), as the probability that a κ -exposed user will adopt the meme before getting $(\kappa + 1)$ -exposed to it, namely

$$P_a(\kappa) = N_a(\kappa)/N_e(\kappa), \quad (1)$$

where $N_e(\kappa)$ is the number of users who were κ -exposed to the information, and $N_a(\kappa)$ is the number of users that were κ -exposed and adopted the meme before becoming $(\kappa + 1)$ -exposed. In other words, $N_e(\kappa)$ and $N_a(\kappa)$ are the numbers of exposure events and adoption events with κ exposers, respectively.⁵

⁵To decrease the impact of noise, we sum up the numbers of events over all memes and then compute $P_a(\kappa)$. Thus, in comparison with the traditional definition of adoption probability, we weight popular hashtags proportionally more.

We extend the above standard definition of adoption probability to include a topical alignment. Namely, we define the conditional probability that a κ -exposed user will adopt the meme before getting $(\kappa+1)$ -exposed when the topical alignment is S , namely

$$P_a(\kappa, S) = N_a(\kappa, S)/N_e(\kappa, S), \quad (2)$$

where $N_e(\kappa, S)$ and $N_a(\kappa, S)$ are the numbers of exposure events and adoption events with the social exposure κ and the topical alignment S . One can obtain one-dimensional adoption probabilities $P_a(\kappa)$ or $P_a(S)$ by integrating $P_a(\kappa, S)$ over κ or S , respectively, which gives

$$P_a(S) = N_a(S)/N_e(S). \quad (3)$$

The adoption probability versus social exposure $P_a(\kappa)$ has been measured in multiple studies. In contrast, to our knowledge, the adoption probability versus a topical alignment $P_a(S)$ is measured for the first time in this study.

4.2 Diffusion via internal and external sources

Information can spread from sources that are internal or external to a given system (in our case, to Twitter). In general, it is difficult to distinguish between internal and external exposures and to estimate their impact on information diffusion (Bakshy et al. 2012). Even if we know that a user follows a person that tweeted some meme and that this person adopted this meme afterwards, we still do not know whether this adoption was caused by the exposures via Twitter followers or via some other external sources. Furthermore, in the next section we show that external adoptions are frequent for information diffusion on Twitter.

Fortunately, our theory of adoption probability allows us to differentiate between information diffusion via internal and external sources, i.e.,

$$P_a(\kappa^i, \kappa^e, S) = P_a^i(\kappa^i, S) + P_a^e(\kappa^e, S). \quad (4)$$

In this study, the internal exposures are via the follower network of Twitter and we can estimate the number of such exposures. By contrast, exposures via external sources are opaque to us. In other words, we do not know κ^e , so we marginalize it out from the last equation obtaining

$$P_a(\kappa, S) = P_a^i(\kappa, S) + P_a^e(S), \quad (5)$$

where $\kappa \equiv \kappa^i$ and $P_a^e(S) = \int P_a^e(\kappa^e, S)P(\kappa^e|S)d\kappa^e$. Next, we note that, if no internal exposures are present, i.e., $\kappa^i = 0$ and $P_a^i(0, S) = 0$, then the only influence remaining is external, thus

$$P_a(0, S) = P_a^e(S). \quad (6)$$

This simple formula shows how to estimate the probability of adoption via external exposures as a function of topical alignment. Combining Equations 5 and 6 allows us to estimate the *joint adoption probability via internal exposures*, namely

$$P_a^i(\kappa, S) = P_a(\kappa, S) - P_a(0, S). \quad (7)$$

Finally, we compute the marginalized internal adoption probabilities $P_a^i(\kappa)$ and $P_a^i(S)$, as follows:

$$P_a^i(S) = P_a(S) - P_a(0, S), \quad (8)$$

$$P_a^i(\kappa) = P_a(\kappa) - \int P_a(0, S)P(S|\kappa)dS. \quad (9)$$

These two formulas subtract the probability of adoption caused by external sources from the probability of any adoptions. Interestingly, Equation 9 subtracts the expected value of adoption probability given that the conditional probability distribution of similarity is $P(S|\kappa)$. Intuitively, even if an individual is exposed to a meme on Twitter via $\kappa > 0$ internal social exposures, he may adopt the meme from another topically related external source; we estimate and subtract the impact of such external phenomena on adoption probability via the term $\int P_a(0, S)P(S|\kappa)dS$.

4.3 Persistence of spreading

We define the persistence of information diffusion as the ability of multiple exposures to boost information adoption. We introduce a new metric of spreading *persistence* (Romero, Meeder, and Kleinberg 2011) to capture this complex contagion effect. We define this persistence as the ratio between the average internal adoption probability at several exposures and the adoption probability at one exposure, namely Persistence = $\langle P_a^i(\kappa) \rangle_{\kappa>1} / P_a^i(\kappa = 1)$, where $\langle \cdot \rangle$ is an average. Intuitively, our persistence metric captures if multiple exposures boost information diffusion in comparison with a single exposure.

5 Impact of topicality on adoption probability

Here, we characterize topical and non-topical mechanisms of information diffusion by studying the adoption probability versus social exposure and topical alignment. We analyze separately the probability of adoption caused by internal and external exposures. Among others, we show that the laws of topical information diffusion apply both to internal and external information diffusion.

5.1 Topicality of memes vs adoption probability

We begin by measuring the distribution of social exposure κ and topical alignment $S(\text{user, meme})$ for adoption and exposure events of topical and non-topical memes (Figure 3). Prior works show that information cascades tend to be shallow and have multiple seed adopters who learn about the memes not from their friends, but from some external sources (Goel, Watts, and Goldstein 2012). Here, we confirm that over 30% of hashtag adoptions and over 40% of URL adoptions happened without a direct exposure through follower links (see y-value at $\kappa = 0$ in Figures 3A-B). Furthermore, we notice that the distribution of $S(\text{user, meme})$ differs for topical and non-topical memes. In particular, the distributions of topical alignment $S(\text{user, meme})$ are skewed to higher values for topical memes than non-topical memes. For instance, about 40% of topical hashtags get adopted when the topical alignment is higher than 0.4, whereas only 2% of non-topical hashtags are adopted when the alignment is higher than 0.4 (compare blue and red solid lines in Figure 3C). Similar difference is present between topical and non-topical URLs (Figure 3D). This signifies that topical memes spread differently than non-topical memes.

Also, there exists a large difference between the cumulative distributions of adoption and exposure events for topical

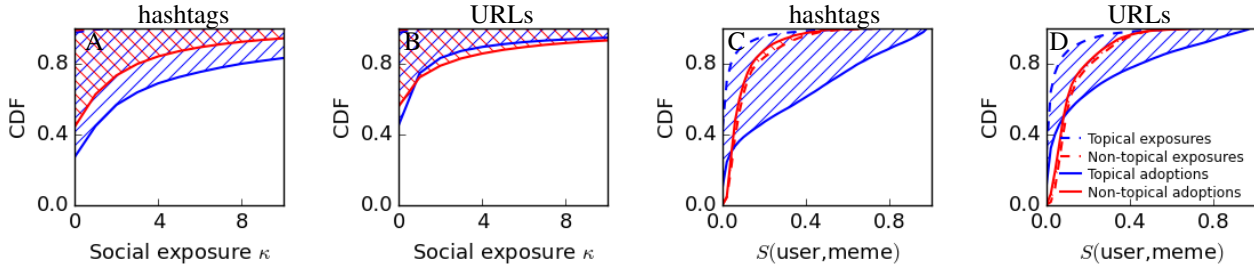


Figure 3: The cumulative distribution of: (A-B) social exposure κ and (C-D) topical alignment $S(\text{user}, \text{meme})$ of adoption events (solid line) and exposure events (dashed line) of hashtags and URLs. The blue lines correspond to topical memes, whereas the red lines correspond to non-topical memes. The filled area marks the difference between the distributions for adoption and exposure events.

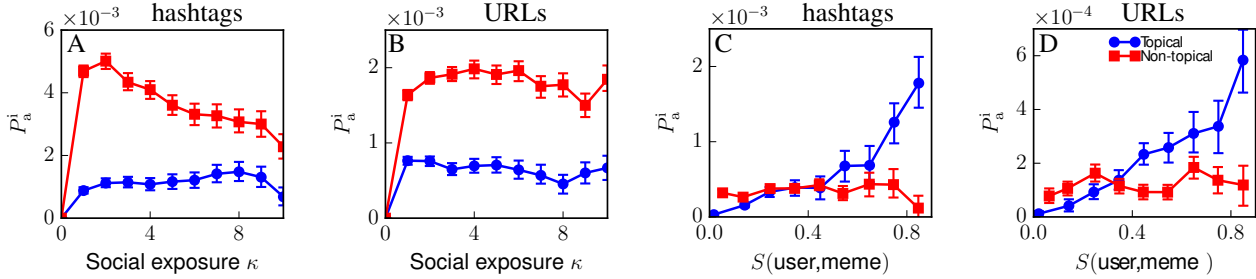


Figure 4: The adoption probability of hashtags and URLs as a function of: (A-B) social exposure κ , (C-D) topical alignment $S(\text{user}, \text{meme})$ for topical (blue circles) and non-topical (red squares) memes. The error bars mark 95% confidence intervals (from BCA bootstrap).

memes (compare solid and dashed blue lines in Figure 3). This difference can be quantified with Kolmogorov-Smirnov distance between the cumulative distributions, which is equivalent to the maximal vertical width of the corresponding filled area in Figure 3. Because the K-S distance is high for topical memes both for social exposure κ and topical alignment $S(\text{user}, \text{meme})$, one can expect that both these properties can act as important discriminators for predicting whether an exposure event will lead to adoption or not. By contrast, for non-topical memes the K-S distance between the distributions of the topical alignment is almost zero (red lines in Figures 3C-D). This signifies that the topical alignment cannot predict the adoption of non-topical memes.

Next, we measure the probability of adoption via internal sources $P_a^i(\kappa)$, finding that it sharply increases with social exposure, reaching a maximum for hashtags and a maximum or a plateau for URLs (Figures 4A-B). The probability decreases for higher number of exposures, because a user who does not adopt a meme after being exposed to it once or twice, will likely not adopt it, even if she is exposed to it by other people. Interestingly, the adoption probability is at least twice higher for non-topical memes than for topical memes. This indicates that non-topical memes are more viral and reach larger audiences than topical memes (Weng and Menczer 2014). In fact, non-topical hashtags have on average 934 adopters and topical hashtags have 500 adopters (significant difference, Mann-Whitney U test, $p = 0.04$). In other words, non-topical memes spread broadly, while topi-

cal memes diffuse more narrowly.

Finally, we measure the adoption probability as a function of the topical alignment between user's interests and the meme, namely $P_a^i(S(\text{user}, \text{meme}))$. We find that the adoption probability curve has different shapes depending on meme's topicality (Figures 4C-D). For topical memes, the adoption probability increases remarkably steadily with the topical alignment, both for hashtags and URLs. In stark contrast, the adoption probability remains flat for non-topical memes. To summarize, topical memes tend to diffuse narrowly among users who are interested in them, while non-topical memes spread broadly among all users.

5.2 Topicality of users vs adoption probability

Here, we describe the diffusion among users with topically narrow and topically diverse interests, i.e., among users with topical and non-topical interests (Figures 5 and 6). Overall, we observe the same tendencies as for topical and non-topical memes. However, these tendencies are less pronounced for topical and non-topical users. For instance, the adoption probability $P_a^i(S(\text{user}, \text{meme}))$ is significantly higher for topical users than for non-topical users (Figures 6C-D); however, here the adoption probability for non-topical users is not flat, but instead increases gradually. Thus, in the remainder of this study, we focus on the results for topical and non-topical memes, because they show a cleaner distinction between the characteristics of topical and non-topical mechanisms of information diffusion.

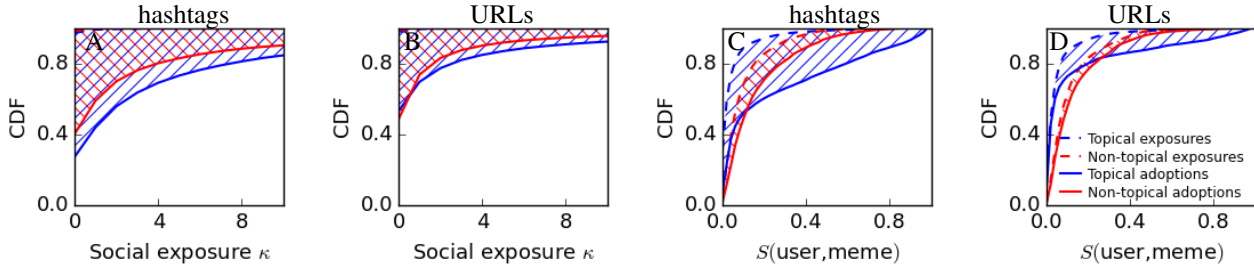


Figure 5: The cumulative distribution of: (A-B) social exposure κ and (C-D) topical alignment $S(\text{user}, \text{meme})$ of adoption events (line) and exposure events (dashed line) of hashtags and URLs. The blue lines correspond to topical users, whereas the red lines correspond to non-topical users.

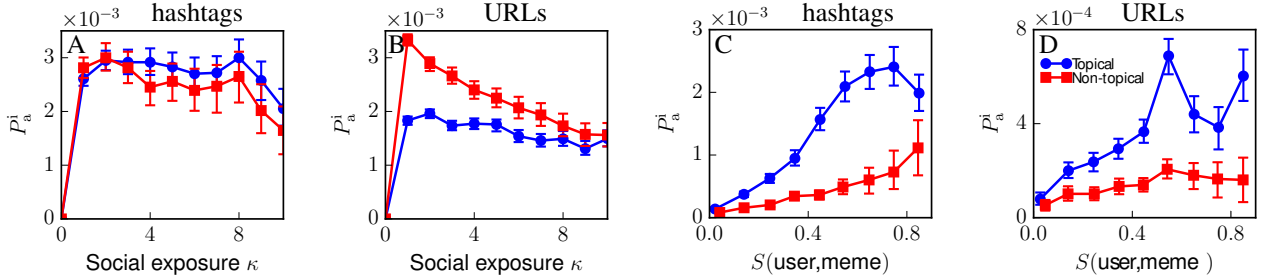


Figure 6: The adoption probability of hashtags and URLs as a function of: (A-B) social exposure κ , (C-D) topical alignment $S(\text{user}, \text{meme})$ for topical (blue circles) and non-topical (red squares) users.

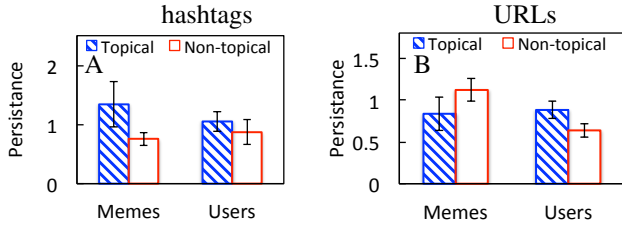


Figure 7: Spreading persistence measured separately for topical and non-topical memes, as well as for topical and non-topical users. The errors bars are 95% confidence intervals.

Finally, we note that multiple social exposures to a meme boost topical diffusion more than non-topical diffusion. To this end, the adoption probability $P_a^i(\kappa)$ drops with social exposures more slowly for topical users than for non-topical users (Figures 6A-B). This is much more evident for the spread of topical hashtags, for which the adoption probability increases with social exposure and reaches a maximum at eight exposures, much more than two social exposures needed to reach the maximum for non-topical hashtags (Figure 4A). To our knowledge, it has not yet been observed that the maximum of adoption probability (aggregating over various hashtags) can be reached at such a high number of social exposures. Using our definition of spreading persistence, we find that the diffusion of topical hashtags is significantly boosted by multiple exposures (compare left bars of Figure 7A). One could expect this result, given the fact

that many of topical hashtags contain novel, unproven, or controversial meaning. Furthermore, we find that the spread of URLs among users with topical interests is also significantly enhanced by multiple exposures (compare right bars of Figure 7B). We note here that the other two cases are insignificant due to the noise in the measurements of the adoption probability (evident in the Figure 6D). It is interesting to see that the complex contagion effect is present not only for topical memes, but also for users with topical interests. A plausible interpretation is that experts or enthusiasts of a given topical domain who decide to spread a certain information risk their topical reputation by doing so. Thus, they are more likely to diffuse information if multiple other people in their social network did so already.

5.3 Seed adopters and the probability of external adoption

Earlier in this section, we noted that a large fraction of users adopt memes without being exposed to these memes by any of their followers. We call such users *seed adopters*. Seed adopters are the initiators of cascades propagating memes in a social network. Importantly, seed adopters at the moment of adoption are not influenced by social exposures internal to Twitter, but instead are influenced by external exposures.

We measure the adoption probability for seed adopters (described by Equation 6) by accounting only for adoption and exposure events at $\kappa = 0$. As before, we split memes

into topical and non-topical. Essentially, we obtain the same shapes of the adoption probability curve for seed adopters as for all adopters (compare Figures 8A-B against Figures 4C-D). For topical memes the adoption probability increases with the topical alignment, whereas for non-topical memes it stays nearly constant. We present two explanations of this finding. First, a seed adopter may have stumbled on the meme through her sources external to Twitter, which likely are also homophilous and topically similar to that adopter. Second, the user may have actively searched for the memes that are topically aligned with her interests via search or by browsing profiles of other users or co-appearing hashtags. Note that this type of exposure can be also considered as external, because it happens via methods external to the follower network of Twitter; essentially, this is how we define external exposures in this study. Finally, the comparison of figures reveals that the diffusion of URLs is nearly twice more intense via external sources rather than internal sources (the probability is twice higher in Figure 8B than in Figure 4D). By contrast, the diffusion of hashtags has similar intensity for external and internal sources. It is easy to interpret this result, knowing that URLs have the global scope of the Internet, whereas hashtags are specific to Twitter.⁶ Because of this distinction, the diffusion of URLs is about twice more intense via external exposures than internal exposures. Note that the diffusion of hashtags via “external” sources may correspond to the exposures that happen within Twitter via ways different than the follower network.

Furthermore, we find that in the domain of topical information diffusion seed adopters are relatively more topically aligned with the adopted memes than other adopters. To compute the relative alignment we divide the alignment averaged over adoption events by the alignment averaged over exposure events, to account for homophily, namely $\langle S(\text{user}, \text{meme}) \rangle_a / \langle S(\text{user}, \text{meme}) \rangle_e$. We find that the relative topical alignment is significantly higher for seed adopters than for other adopters, both for topical hashtags and topical URLs (Figure 9). Intuitively, seed adopters are more interested in the topic of a meme than other adopters, similar to early adopters being more enthusiastic about an innovation than other people (Rogers 1962). This sheds new light on the reason why memes have multiple seed adopters (Goel, Watts, and Goldstein 2012). Naturally, this effect is not present for non-topical memes.

6 Discussion

The findings presented in this study create the following schematic of topical information diffusion in Twitter. Topical memes start to spread from the enthusiasts or experts of the corresponding topic. Then, they diffuse narrowly to other users whose interests are topically aligned with the topics of the memes. Multiple exposures to the same meme tend to boost topical information diffusion, as predicted by the theory of complex contagion, since many topical memes contain novel or controversial information. Restarts of informa-

⁶Twitter is the first platform where hashtags emerged, initially as a social convention in 2007, and then were implemented by Twitter as a part of the system in 2009.

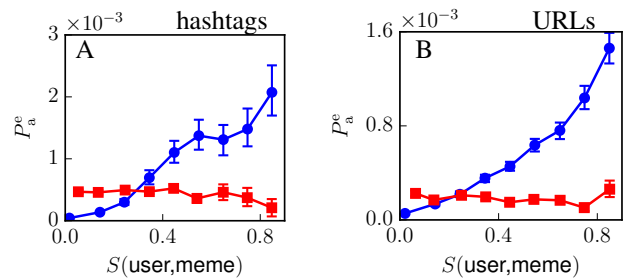


Figure 8: The probability of external adoption ($\kappa = 0$) as a function of topical alignment between a user and a meme for topical (blue) and non-topical (red) memes. The error bars mark 95% confidence intervals.

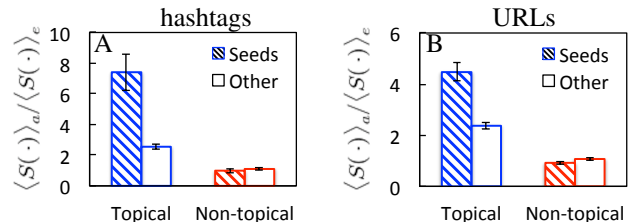


Figure 9: The relative topical alignment for seed and other adopters. The error bars are 99% confidence intervals.

tion cascades are frequent and are initiated by seed adopters who tend to be relatively more topically aligned with the given information than other adopters. Overall, we have shown that topical information diffusion differs from non-topical information diffusion in several meaningful ways.

Our work explores the gap between the studies of information diffusion and recommender systems by showing that users even without personalized recommendations adopt information that is topically related to them. Clearly, recommender systems may strengthen or weaken the natural human tendency to spread topically aligned information. How recommender system affect information diffusion? Could we define a recommender system that facilitates information diffusion without causing a filtering bubble? This is a set of questions that could be addressed in future studies. Furthermore, in this study, rather than proposing and training a supervised model predicting information diffusion, we use a state-of-the-art unsupervised topic modeling method and analyze the impact of topicality on information diffusion independently. As such, our results provide an empirical lower estimate of the impact of topicality on diffusion. Future studies could propose a topic model merged with a diffusion model as an effective recommender system.

To conclude, we note that information diffusion is a fundamental building block of many collective processes happening in social networks, e.g., viral marketing, spreading of political movements and innovation, and opinion formation. Thus, our findings have wide implications for understanding, predicting, and controlling such processes.

Acknowledgments

We thank Komal Agrawal for her initial work on this project, Robert West for his valuable suggestions, Vlad Niculae for his insights into the classification problem, and Vlad Niculae, Manuel Gomez-Rodriguez, and Juhi Kulshrestha for their feedback on early versions of this study. Finally, we thank the anonymous reviewers for helping to improve this work. Also, we acknowledge the support of project IMPECS “Understanding, Leveraging and Deploying Online Social Networks”.

References

- [Backstrom et al. 2006] Backstrom, L.; Huttenlocher, D.; Kleinberg, J.; and Lan, X. 2006. Group formation in large social networks. In *Proc. 12th ACM SIGKDD Int. Conf. Knowl. Discov. data Min. - KDD '06*, 44. New York, New York, USA: ACM Press.
- [Bakshy et al. 2012] Bakshy, E.; Rosenn, I.; Marlow, C.; and Adamic, L. 2012. The role of social networks in information diffusion. In *Proc. 21st Int. Conf. World Wide Web - WWW '12*, 519. New York, New York, USA: ACM Press.
- [Barbieri, Bonchi, and Manco 2012] Barbieri, N.; Bonchi, F.; and Manco, G. 2012. Topic-Aware Social Influence Propagation Models. *2012 IEEE 12th Int. Conf. Data Min.* 81–90.
- [Bi et al. 2014] Bi, B.; Tian, Y.; Sismanis, Y.; Balmin, A.; and Cho, J. 2014. Scalable topic-specific influence analysis on microblogs. *Proc. 7th ACM Int. Conf. Web search data Min. - WSDM '14* 513–522.
- [Blei, Ng, and Jordan 2003] Blei, D.; Ng, A.; and Jordan, M. 2003. Latent Dirichlet Allocation. *J. Mach. Learn. Res.*
- [Burgner, Jamieson, and Blackwell 2006] Burgner, D.; Jamieson, S. E.; and Blackwell, J. M. 2006. Genetic susceptibility to infectious diseases: big is beautiful, but will bigger be even better? *Lancet Infect. Dis.* 6(10):653–663.
- [Centola and Macy 2007] Centola, D., and Macy, M. 2007. Complex Contagions and the Weakness of Long Ties. *Am. J. Sociol.* 113(3):702–734.
- [Centola 2010] Centola, D. 2010. The Spread of Behavior in an Online Social Network Experiment. *Science (80-.)*. 329(5996):1194–1197.
- [Cha et al. 2010] Cha, M.; Haddadi, H.; Benevenuto, F.; and Gummadi, K. P. 2010. Measuring user influence in twitter: The million follower fallacy. In *Proc. ICWSM*, 10–17.
- [Cosley, Huttenlocher, and Kleinberg 2010] Cosley, D.; Huttenlocher, D.; and Kleinberg, J. 2010. Sequential Influence Models in Social Networks. *ICWSM '10 Proc. 4th Int. AAAI Conf. Weblogs Soc. Media* 26–33.
- [Crandall et al. 2008] Crandall, D.; Cosley, D.; Huttenlocher, D.; Kleinberg, J.; and Suri, S. 2008. Feedback effects between similarity and social influence in online communities. *Proceeding 14th ACM SIGKDD Int. Conf. Knowl. Discov. data Min. - KDD 08* 160.
- [Du et al. 2013] Du, N.; Song, L.; Woo, H.; and Zha, H. 2013. Uncover Topic-Sensitive Information Diffusion Networks. In *AISTATS'13*, volume 31, 229–237.
- [Goel, Watts, and Goldstein 2012] Goel, S.; Watts, D. J.; and Goldstein, D. G. 2012. The structure of online diffusion networks. In *Proc. 13th ACM Conf. Electron. Commer. - EC '12*, volume 1, 623. New York, New York, USA: ACM Press.
- [Gomez-Rodriguez, Gummadi, and Schölkopf 2014] Gomez-Rodriguez, M.; Gummadi, K.; and Schölkopf, B. 2014. Quantifying Information Overload in Social Media and its Impact on Social Contagions. In *ICWSM '14 Proc. 8th Int. AAAI Conf. Weblogs Soc. Media*.
- [Granovetter 1978] Granovetter, M. 1978. Threshold Models of Collective Behavior. *Am. J. Sociol.* 83(6):1420–1443.
- [Holley and Liggett 1975] Holley, R. A., and Liggett, T. M. 1975. Ergodic theorems for weakly interacting infinite systems and the voter model. *Ann. Probab.* 643–663.
- [Kempe, Kleinberg, and Tardos 2003] Kempe, D.; Kleinberg, J.; and Tardos, É. 2003. Maximizing the spread of influence through a social network. In *Proc. ninth ACM SIGKDD Int. Conf. Knowl. Discov. data Min. - KDD '03*, 137. New York, New York, USA: ACM Press.
- [Kosinski, Stillwell, and Graepel 2013] Kosinski, M.; Stillwell, D.; and Graepel, T. 2013. Private traits and attributes are predictable from digital records of human behavior. *Proc. Natl. Acad. Sci.* 2–5.
- [Leskovec, Adamic, and Huberman 2007] Leskovec, J.; Adamic, L. A.; and Huberman, B. B. A. 2007. The Dynamics of Viral Marketing. *ACM Trans. Web* 1(1).
- [Ma, Sun, and Cong 2013] Ma, Z.; Sun, A.; and Cong, G. 2013. On predicting the popularity of newly emerging hashtags in Twitter. *J. Am. Soc. Inf. Sci. Technol.* 64(7):1399–1410.
- [McCallum 2002] McCallum, A. K. 2002. MALLETT: A Machine Learning for Language Toolkit.
- [Nakatani 2012] Nakatani, S. 2012. Language Detection with Infinity Gram.
- [Owoputi et al. 2013] Owoputi, O.; O'Connor, B.; Dyer, C.; Gimpel, K.; Schneider, N.; and Smith, N. A. 2013. Improved Part-of-Speech Tagging for Online Conversational Text with Word Clusters. In *Proc. NAACL-HLT 2013*.
- [Ricci, Rokach, and Shapira 2011] Ricci, F.; Rokach, L.; and Shapira, B. 2011. *Introduction to Recommender Systems Handbook*. Boston, MA: Springer US.
- [Rogers 1962] Rogers, E. M. 1962. *Diffusion of innovations*. Simon and Schuster.
- [Romero, Meeder, and Kleinberg 2011] Romero, D. M.; Meeder, B.; and Kleinberg, J. 2011. Differences in the mechanics of information diffusion across topics. In *Proc. 20th Int. Conf. World wide web - WWW '11*, 695. New York, New York, USA: ACM Press.
- [Strang and Soule 1998] Strang, D., and Soule, S. A. 1998. Diffusion in organizations and social movements: From hybrid corn to poison pills. *Annu. Rev. Sociol.* 265–290.
- [Weng and Menczer 2014] Weng, L., and Menczer, F. 2014. Topicality and Social Impact: Diverse Messages but Focused Messengers. 9.

[Weng et al. 2010] Weng, J.; Lim, E.-p.; Jiang, J.; and He, Q. 2010. TwitterRank: finding topic-sensitive influential twitterers. In *Proc. third ACM Int. Conf. Web search data Min. - WSDM '10*, 261. New York, New York, USA: ACM Press.

[Yang et al. 2012] Yang, L.; Sun, T.; Zhang, M.; and Mei, Q. 2012. We Know What @You #Tag: Does the Dual Role Affect Hashtag Adoption? In *Proc. 21st Int. Conf. World Wide Web, WWW '12*, 261–270. New York, NY, USA: ACM.

[Yuan et al. 2016] Yuan, N. J.; Zhong, Y.; Zhang, F.; Xie, X.; Lin, C.-y.; and Rui, Y. 2016. Who Will Reply to/Retweet This Tweet?: The Dynamics of Intimacy from Online Social Interactions. In *Proc. Ninth ACM Int. Conf. Web Search Data Min. - WSDM '16*, 3–12. New York, New York, USA: ACM Press.