

Fragile Online Relationship: A First Look at Unfollow Dynamics in Twitter

Haewoon Kwak, Hyunwoo Chun, and Sue Moon

Department of Computer Science, KAIST
291 Daehak-ro, Yuseong-gu, Daejeon, Korea
{haewoon, hyunwoo}@an.kaist.ac.kr, sbmoon@kaist.edu

ABSTRACT

We analyze the dynamics of the behavior known as ‘unfollow’ in Twitter. We collected daily snapshots of the online relationships of 1.2 million Korean-speaking users for 51 days as well as all of their tweets. We found that Twitter users frequently unfollow. We then discover the major factors, including the reciprocity of the relationships, the duration of a relationship, the followees’ informativeness, and the overlap of the relationships, which affect the decision to unfollow. We conduct interview with 22 Korean respondents to supplement the quantitative results. They unfollowed those who left many tweets within a short time, created tweets about uninteresting topics, or tweeted about the mundane details of their lives. To the best of our knowledge, this work is the first systematic study of the unfollow behavior in Twitter.

Author Keywords

Unfollow, computer-mediated communication, Twitter, online relationship

ACM Classification Keywords

H.5.0 Information Systems: Information Interfaces and Presentation - General

General Terms

Human Factors

INTRODUCTION

Relationship formation and dissolution are two basic processes of relationship change and evolution in personal networks. Studies of relationship formation and dissolution mostly rely on surveys and interviews, both of which require considerable effort in terms of time and labour. Online social networks (OSN) aid researchers in at least two ways, such as (i) they contain a huge archive of human behavior related to online relationships, and (ii) they allow easy access. Studies of online relationship formation are straightforward, as most OSNs offer simple means of establishing

online relationships, often referred to as a ‘friend’. By contrast, research on the topic of online relationship dissolution has not been extensively conducted due to the lack of data; an online friend relationship remains rigid regardless of the actual relationship [28]. Researchers thus use proxies to represent the state of relationship dissolution. For example, a study of relationship dissolution in email networks assumes that the disappearance of online activities (the exchange of emails) reflects this type of dissolution [16]. However, a disappearance of communication cannot be directly translated to the dissolution of a relationship in most cases, because it is difficult to capture all communication means [23] and to regard the absence of an event as strictly intentional. The key insight behind this work is that unfollow in Twitter represents a user’s explicit expression of the dissolution of an online relationship and breaks off the non-reciprocal online relationship.

Twitter has been redefining human behavior logging on the strength of worldwide popularity, brief text messages, and well-supported application programming interface (API) to access. Twitter allows a user to establish a one-way relationship known as *follow*. This is one of the most distinguishing features of Twitter, in contrast to the reciprocal friendship in most other OSNs, e.g., friends in Facebook. Literally, neither an invitation nor an acceptance is required for following; people can freely and easily follow others. When a user logs into Twitter, s/he sees tweets from those s/he follows (*followees*). The stream of followees’ tweets is called a *timeline*. A user can easily stop following (*unfollow*) and needs no confirmation from the followee to do so. Unfollow, thus, is not a proxy but a verifiable action of breaking an online relationship. In the rest of the paper, we use unfollow both as a noun and a verb.

In this work, we analyze the dynamics of the unfollow behavior to understand online relationship dissolution. The two research questions explored here are: (i) what are the characteristics of the unfollow behavior? and (ii) why do people unfollow others? To address the first research question, we collected daily snapshots of the follow relationships of 1.2 million Korean-speaking users over the course of 51 days as well as their tweets. By comparing the daily snapshots, we confirm that unfollow is prevalent in Twitter. We have found that the reciprocity of the relationship, the duration of the relationships, the followees’ informativeness, and the overlap of relationships are critical in the decision to un-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2011, May 7–12, 2011, Vancouver, BC, Canada.

Copyright 2011 ACM 978-1-4503-0267-8/11/05...\$10.00.

follow. For the second research question, we conducted interviews with 22 users to determine their motivations behind the unfollow behavior. Our survey respondents answered that they unfollowed those who left many tweets within a short time, left tweets about topics uninteresting to them, or left tweets about mundane details of daily life. To the best of our knowledge, this work is the first systematic look at the unfollow behavior in Twitter.

For the rest of the paper, we first review the literature on friendship dissolution problems and state-of-the-art Twitter analyses. We then explain research questions and give the outline of this work. We describe our data collection methodology of crawling Twitter and interviewing Twitter users. We then demonstrate interesting unfollow patterns and factors that impact the decision to unfollow from large-scale data. Next, we present qualitative analyses of the underlying user motivation behind unfollow according to interviews with the Twitter users. We finally conclude by discussing implications for future research.

PREVIOUS LITERATURE

For better understanding this work, we review the literature on three different areas. We begin with the research on the topic of off-line relationship dissolution. Next, we introduce studies of the characteristics of user behavior in Twitter. Finally, we review models and empirical studies of the link dynamics in networks.

Friendship Dissolution

A few longitudinal studies have reported on off-line relationship dissolution. Mollenhorst have found that a half of adult friendships change over the period of seven years [25]. Two studies have observed that social networks among younger children have a higher retention rate [6, 12]. Homophily plays a decisive role in off-line relationship retention; several studies report that friendship is more stable when people are more alike [16, 30]. Homophily is a combined process of establishing friendships with those who are like you and acting like your friends [15]. Our interviewees confirm that they often stop following those who write tweets about topics uninteresting to them.

User Behavior Analysis in Twitter

Pioneering work by Java *et al.* has sampled 76,000 Twitter users and found users to be clustered based on topics [14]. Kwak *et al.* have crawled the near-complete Twittersphere and collected profiles of 41.7 million users and 1.47 billion tweets of trending topics. They reported that Twitter carries the characteristics of a news medium rather than works as a typical online social network [18]. Other quantitative analyses of user behavior in Twitter also looked at the perspective of information delivery: homophily in terms of the shared interests [32], informational communication at work [33], and information propagation [20]. Facebook has been the focus of a good number of qualitative user behavior analysis [9, 19, 28], but few have covered qualitative analysis of Twitter yet. We study the unfollow behavior from the large-scale data and supplement with qualitative interviews for clear understanding of motivations behind unfollow.

Link Dynamics

From theoretical models to empirical studies of link dynamics in networks, link removal has been largely ignored due to lack of verifying data. The link prediction problem is “given a snapshot of a social network at time t , we seek to accurately predict the edges that will be added to the network during the interval from time t to a given future time t' ” [22]. It deals only with newly added links but not removed links. The generative models in [3, 21] account only for link additions in the network evolution, and so are previous empirical studies [10, 17]. A recent study focuses on the structural predictors of tie formation particularly in Twitter [11]. Our study reports on the prevalence of link removal and major factors to affect link removal in the online relationship network of Korean Twitter users. It can aid researchers in creating new models and verifying the extensibility of conventional models.

THE GOAL OF THIS PAPER

Our goals are to observe the characteristics of the unfollow behavior and to understand why people unfollow.

R1: What are the characteristics of the unfollow behavior?

We reveal unfollow patterns through quantitative data analysis. We begin with the frequency of unfollow. We expect that unfollow is prevalent because: (i) follow relationships do not need to be reciprocal, (ii) Twitter does not notify unfollow to those who are unfollowed, and (iii) unfollow is done by just one click. We also find temporal correlation in the unfollow behavior. We then demonstrate the correlation between the frequency of interaction and the likelihood of unfollow. This shows the validity of general assumptions that the less frequency of interaction could be a sign of the dissolution of an online relationship.

Next, we analyze the stabilization of the relationships; do people tend to unfollow new followees or old ones? If older relationships are more broken, people continuously want the freshness in their timelines. Or if newer relationships are more broken, older followees mean that people are satisfied with them. If no correlation is observed, people unfollow a followee no matter how long a relationship has lasted. We also aim to find evidences that people tend to control the number of followees.

R2: Why do people unfollow?

Beyond the observations of the unfollow behavior, we answer what factors affect unfollow. We examine the motivations behind unfollow through not only the analysis of huge data but also user interviews.

First is the informativeness of a followee. Previous work highlighted the perspective of information delivery in Twitter. If the purpose of following is mostly subscribing to tweets, more informative followees could be hardly unfollowed.

Next, we study how the reciprocity of the relationships affects unfollow. We compare the likelihood of unfollow in

reciprocal relationships to that in unreciprocated relationships. We also find the correlation between the overlap of relationships of two users and the likelihood of unfollow.

Then, we conduct interviews with 22 Korean users. We divide our questions into four categories: the motivation behind unfollow, the awareness of being unfollowed, the marketing campaigns driven by company accounts, and the ranking of followees in terms of the likelihood of unfollow.

DATA COLLECTION

The main challenge is to obtain traces of the unfollow behavior because Twitter does not offer the explicit records of unfollow. We solve this problem by comparing daily snapshots of each user’s follow relationships and detecting any disappearance of a follow relationship as unfollow.

As of September 2010, the total number of Twitter users is close to two hundred million, too large for a few tens of crawlers to collect the entire user space within a short period of time. We had to choose a sampling strategy. Instead of choosing a random sampling methodology, we wanted a sample with a cohesive cultural and societal bound so that we could compare personal interviews in the same context. A typical choice in previous studies of online relationships would have been a university [19, 27, 29]. Thanks to the open API of Twitter, we could expand our sample from a university to a country bounded by the language.

We have chosen the exhaustive set of Korean-speaking users for the following reasons. First, it is like a miniature of the entire Korean society. Almost all walks of life are represented, such as politicians, sports stars, TV anchors, writers, students, housewives, labor unionists, etc., and relationships among them, such as celebrities and fans, political parties and their followers, and manufacturers and consumers. This offers relatively rich user behavior, in contrast to the limited scope of college life in previous studies. Further, we can get a longitudinal data with a fine-grained time resolution, as the pool of all Koreans is a manageable size to crawl in a single day. The other reason is authors are familiar with the culture of the online Korean community. We can grasp the context of tweets about almost every online or off-line issues in the Korean community.

We have collected the followers and followees of every Korean Twitter user. Any Twitter user who fell to one or more of the following four conditions was deemed Korean: (i) one wrote a tweet in Korean; (ii) one’s biographical information was written in Korean; (iii) one’s location was written in Korean; and (iv) one’s screen name was written in Korean. We made daily snapshots of the entire Korean users from June 25th, 2010 to July 15th, 2010 and from August 2nd to August 31, 2010. We could not collect snapshots from July 16th to August 1st due to technical problems. In the rest of the paper we label social graphs from two isolated periods as $G(I)$ and $G(II)$, respectively.

Next, we have crawled all tweets written by Korean users until the last day of $G(II)$, August 31st, 2010. Through the col-

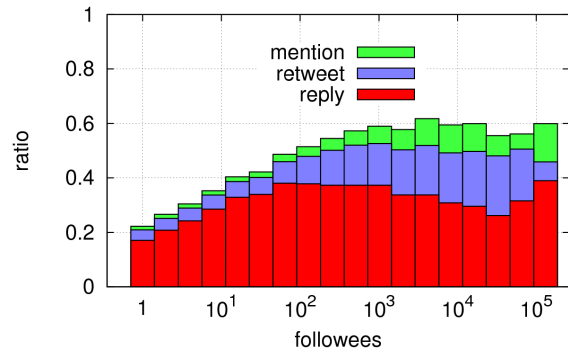


Figure 1. Proportion of tweets of each type

lected tweets, we can reconstruct who interacted with whom by a reply or a mention, and who acknowledged whom by a retweet or a favorite. The retweet has been typically considered as an indicator of the tweet’s quality [18, 20], and we consider a ‘favorite’ as such an indicator as well.

We note that the social and cultural norms about friendship, commitment, and online services might play out differently in Korea than in other countries. We thus need more data to show the generalizeability of our findings to non-Korean users. We plan to expand our dataset across both time and space for external validity in future work.

BASIC STATISTICS OF THE KOREAN TWITTER NETWORK

We first examine the basic statistics of Korean Twitter users. The number of Korean users in the last day of $G(I)$ is 870,057 and that of $G(II)$ is 1,203,196. The average number of new incoming users per day is 7,599 in $G(I)$ and 8,515 in $G(II)$.

We define the reciprocity of a network as the ratio of user pairs of mutual relationships over the total number of user pairs of either one-way or mutual relationships. Kwak *et al.* have reported a very low reciprocity of 22.1% in the near-complete Twitter network [18]. We find a quite high reciprocity among Korean users: in $G(I)$ we see 56-58% of reciprocity and in $G(II)$ 61-62%. It is still lower than the reciprocity observed in other services, such as 84% of Yahoo! 360 [17]. Interestingly, the reciprocity has increased slightly over time. We additionally discover that the average number of followees increases from 59.7 to 75.7 during the 51 days. These two observations illustrate a densification process in a directed network as similar as that reported in an undirected network [21].

A tweet can be classified into four categories by its purpose: a regular tweet, a reply, a mention, or a retweet. A reply and a mention are conversational, while a retweet is informational. How does a user utilize the four categories vary depending on the number of followees? Figure 1 shows the average proportion of tweets of each type written by the users with the same number of followees. We omit the proportion of regular tweets in the figure. We read from the figure that the percentage of replies increases linearly to the




Figure 2. Number of relationship formation and dissolution per day

number of followees up to 100, signifying that the more one follows, the more interactive one becomes. Once the number of followees exceeds 100, then the proportion of replies stays constant or decreases slightly, but that of retweets grows (the middle section in a bar expands from $x = 10^2$ to $x = 10^4$ and a little beyond). That is, the more followees one has, the more informational one becomes rather than conversational. We note that mentions are by far the smallest in ratio. Twitter offers buttons for a reply and a retweet, but none for a mention. The lack of a convenience feature can result in low usage of mentions, and the addition of the feature might increase the ratio in future.

QUANTITATIVE ANALYSIS OF UNFOLLOW BEHAVIOR

We demonstrate that the unfollow behavior is prevalent in Twitter and the low frequency of interaction cannot signify unfollows due to the passivity in most relationships. Especially, the latter shows the difficulty that low frequency of interaction is regarded as online relationship dissolution in Twitter. Then we discover some factors affecting the unfollow behavior, such as the mutuality of relationships, the duration of a relationship, the followees' informativeness, and the overlap of the relationships of two users.

Pervasive Unfollow

We begin with the network-wide statistics of unfollow. We count the number of relationship formation and dissolution by comparing snapshots from consecutive days and plot them in Figure 2. The lower bars labeled 'removed' represent unfollows and the upper bars are new follows. The new follow is divided into old and new: the former refers to relationship formation among existing users and the latter relationship formation with those users who newly joined on the day. We can see that the frequencies of follows and unfollows are strongly correlated with the Pearson correlation coefficient of 0.715 ($p \ll 0.0001$). Overall, 171,131 users in $G(I)$, 289,444 users in $G(II)$, and 360,321 users in $G(I)+G(II)$ have stopped following at least once. During the 51 days about 30% of users had unfollowed at least once. If we filter out inactive users, who did not establish or break even a single relationship during the 51 days, the ratio increases to 43%. These numbers could increase if we could expand the time window of data collection. The average of unfol-




Figure 3. Number of communication partners vs. number of followees

lows per person is 15.4 in $G(I)$ and 16.1 in $G(II)$. We conclude that unfollow is quite pervasive in Twitter.

Unfollows Frequent and Mostly Singular, But Often Clustered

To understand the temporal characteristics of the unfollow behavior, we analyze how many followees a user unfollow per day and how often a user unfollows. The resolution of our data is a single day for the interval of unfollow. We observe highly skewed unfollow patterns. About 66% of unfollows occurred singularly in a day, while 10% of unfollows occurred with at least other 5 unfollows. In addition, a few users who have many followees unfollow more than a hundred a day. We conjecture that they use some automatic tools to unfollow those who do not follow them back. We also observe that 90% of time intervals between days of unfollows of a user are less than 9 days.

In summary most unfollows are a singular event of the day, while the remaining unfollows appear clustered. Clustered or not, unfollows are a frequent event occurring once in 9 days for 90% of the cases.

Passiveness in Relationships

Unfollow is a much stronger expression of losing interest than diminished interaction between two people. Yet still, is unfollow an extreme form of diminished interaction? To paraphrase the question, can we predict an unfollow from the reduced interaction or not? This illustrates whether a proxy of frequency of interaction for online relationship dissolution is applicable in Twitter.

We first look at how many people a user actually interacts with. We define *communication partners* of a user as those who exchange tweets with the user reciprocally at least once by a mention, a reply, or a retweet. That is, a reply, a mention, or a retweet was sent from one user to the other and vice versa. This is a strong definition of communication in a network with a relatively low reciprocity. Figure 3 shows the correlation between the number of followees a user has and that of communication partners. For clarity of presentation we use log-scale in both x and y axes and log-scale bins for data points. At a first glance the plot exhibits a linear trend, but because the x -axis in log scale, a user is actually inter-

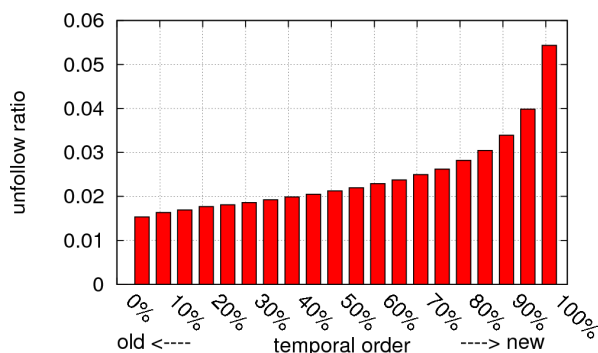


Figure 4. Aggregated number of broken relationships vs. temporal order of relationship establishment in a personal network

acting with an order of magnitude smaller number of people than one’s followees. The number of communication partners is positively correlated with the number of followees, as predicted, but its number is an order of magnitude smaller than that of followees. For example, users who follow 1000 people communicate with only about 70 people on average. This result shows that the majority of relationships do not result in social interaction in Twitter. Interestingly, we find a similar phenomena in other OSNs; Facebook users interact with a relatively small number of their own friends [8] as well as Cyworld users do [7]. These can be explained by the Dunbar’s number, defined as ‘the theoretical cognitive limit to the number of people with whom one can maintain stable social relationships’.

We observe the same level of low activity over most follow relationships. Out of 104, 116, 484 follow relationships, 85.6% of relationships involve no activity at all and 96.3% of relationships involve 3 or fewer. This result reveals the very passive nature of follow relationships. People mostly just subscribe to followees’ tweets, and they do not send a reply, a mention, or a retweet.

We compare the average numbers of sending tweets to those later unfollowed and to staying relationships. We find the difference is less than one tweet for the 99% confidence interval as most relationships are passive and incur no interaction. If we exclude passive relationships in calculation, broken relationships involve less activity than unbroken staying ones: the mean number of sending tweets is 4.107 for the former and 5.850 for the latter. The difference is statistically significant by the two-sample t -test ($p \ll 0.0001$).

Relationship Stabilization

Is a user less likely to unfollow those who the user has been following long enough? In order to gauge the duration of follow and correlate it with the likelihood of unfollow, we need to know when each relationship is established, but Twitter does not offer the information. Instead, Twitter offers the temporal order of the establishment of relationships in the personal network. We thus can recognize the relative order in relationship establishment of followees for each user. Now which relationship is more likely to be broken, old




Figure 5. Number of followees vs. Number of unfollowees

or new? We use 20 bins and put unfollows according to the ego-centric temporal ordering of relationship establishments. Because we choose the bin size of 20, we only consider users with more than 20 followees. For example, let us consider a user with 40 followees unfollowed the very first followee. Then the very first bin in Figure 4 would be incremented by one. Even though the y -axis in the graph is presented in likelihood, it is basically the same. We observe an increasing trend in likelihood of unfollow towards newer followees. The unfollow ratio increases significantly in the rightmost bin, which represents the most recent followees. The figure concludes that the longer the followee remains, the less likely the relationship is broken by unfollow.

No Cap in the Number of Followees

Here we inspect the overhead of managing online relationships based on the the number of followees a user has. We seek to find evidences whether people make an effort to maintain a manageable number of followees. We depict a correlation between the number of followees and that of unfollowees for each user in Figure 5. We do not find any evidence that people try to put a cap on the number of followees. Instead, the ratio of the number of followees to that of unfollowees remains between 1/12 and 1/10.

Reciprocity of Relationships

Reciprocal follow relationships can bring emotional closeness to both users and thus reduce the likelihood of unfollow. We compare the likelihood of unfollow in reciprocated relationships to one-way relationship. The likelihood that one of reciprocated relationships is broken is 0.0529, and that of one-way relationship is 0.1228. The one-way relationship is twice more likely to be broken than the reciprocated relationship. With a two-sample t -test, we see that the difference is statistically significant ($p \ll 0.0001$). We further investigate the likelihood of the remaining relationship to be broken after one of two-way relationships is broken. Interestingly, we observe that the likelihood increases to 0.2345. The remaining relationship is more likely to be broken than likelihood of one-way relationship to be broken. Note that 0.2345 is computed with the data for 51 days, and this likelihood can increase with the longer duration of data.

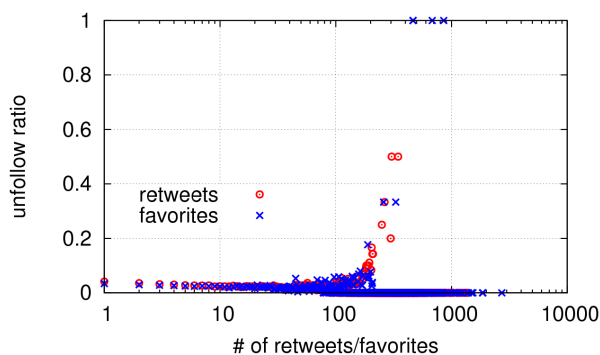


Figure 6. The correlation between followee’s informativeness and the likelihood of unfollowing the followee

Followees’ Informativeness

People can choose to unfollow a followee when they are no longer interested in reading the tweets of the followee. In other words, the informativeness of user’s tweets can be the opportunity cost of unfollowing the user. The most common way to assess the informativeness of user’s tweets is how many times user’s tweets are retweeted [18, 32]. Similarly, it is possible to count how many people mark the user’s tweets as their favorites. A helpful analogy for the favorites in Twitter is the bookmarks in web browsers; a user may mark a tweet as a favorite for future reading.

There is neither a common agreement nor the systematic study on which one of a retweet and a favorite is a better metric to evaluate informativeness, as a retweet and a favorite have different purposes and semantics. People retweet a followee’s tweet to broadcast it to their own followers, whereas marking a favorite mainly for personal use. Here we consider both a retweet and a favorite as an indicator of user’s informativeness and compare two results.

If we aggregate retweets and favorites of a followee by all the followers, we could inaccurately reflect the followee’s informativeness to a user. For example, even if user *A*’s tweets are marked as user *B*’s favorites many times, user *C* may not regard user *A* as informative due to different interests. Therefore, we define informativeness strictly between a user and a followee; the informativeness of my followee to me is measured by only my retweets and favorites.

Figure 6 shows the correlation between a user’s informativeness and the likelihood of being unfollowed. A few unfollow ratios above 0.2 are aberrant data points as only they represent only singular cases. We only focus on the majority of the cases that have the unfollow ratio lower than 0.1. It means that once a relationship involves either retweets or favorites, such relationship is less likely to be broken, while the probability of a relationship randomly picked is 0.114. The data points by the retweets are mostly under those by the favorites below for retweets and favorites under 100. It indicates retweets are a better metric than favorites in measuring informativeness: if one retweets a followee’s tweet, the user is less likely to unfollow the followee than in the case of favoriting a tweet.




Figure 7. The correlation between the overlap of two users’ relationships and the likelihood of the relationship between them to be broken

Shared Relationships

Onnela *et al.* find a positive correlation between the number of common friends of two users and the duration of calls between them in a large undirected mobile phone call network [26]. Similarly, in a directed network, outlink equivalence and inlink equivalence of two nodes are defined as the overlap in outlinks and inlinks of the two nodes, respectively [31]. In the domain of Twitter, the common followees map to outlink equivalence, and the common followers to inlink equivalence.

For every pair of users who have a follow relationship between them, we first compute the outlink equivalence of the pair. Then we normalize it by dividing the number of the union of followees of both users. This is basically the Jaccard coefficient for the two users’ followees. We ignore the two users themselves in set operations. Unless we ignore two users in set operations, the maximum normalized value cannot reach 1 because a user cannot follow oneself. For the input equivalence, we normalize the values in the same manner. Figure 7 shows the correlation between the unfollow ratio of relationships and the normalized input and output equivalences. The unfollow ratio slowly increases up to 0.24 of the normalized outlink equivalence. In the rest of the graph, the unfollow ratio decreases while the normalized input and output equivalences increase. This result suggests that the overlap in followers or followees can be a proxy of tie strength in retaining relationships.

QUALITATIVE ANALYSIS - INTERVIEWS

In order to examine the motivation behind unfollow, we supplement interviews to this work. We recruited 22 survey respondents (11 males and 11 females). The age of study participants ranges from 22 to 36 years. Including the number of participants’ followees, followers, and written tweets, we summarize the demographic breakdown of the 22 respondents in Figure 8.

We began an interview with a common question on their unfollow experience and proceeded with follow-up questions. The interviews lasted between 20 and 30 minutes for face-to-face interviews and between 30 and 40 minutes for Instant messengers, such as Google Talk. The difference in the

	Mean	Median	Min	Max	Std. dev.	Distribution
Age	27.3	27	22	36	3.7	
Favorites	80.7	1	0	851	199.0	
Followers	846.7	164.5	5	8,772	2,053.9	
Followings	600.4	144.5	5	7,103	1,562.7	
Tweets	3,325.8	583.5	5	30,639	7,220.5	
Registered days	449.2	471	14	766	179.1	

Figure 8. Demographic of 22 participants (11 males and 11 females)

interview duration comes mainly from the communication speed in speaking and typing. All face-to-face interviews were recorded by the digital camcorder and transcribed, and instant messenger interviews were recorded as client-side logs. All interviews were conducted in Korean and transcribed later by the first author. We transcribed the interview logs according to standard coding techniques reviewed by Miles and Huberman [24]. As having all tweets written by the participants, we minimized the informant accuracy led by retrospective data [5].

Our first question was why a participant decided to unfollow. Even though we had records of the participant’s unfollow during $G(I)$ and $G(II)$, we did not explicitly mention it in order not to make the participants feel uncomfortable. Most participants easily recalled multiple reasons to unfollow, even without browsing one’s own Twitter page. Our next question was whether the participant thought the unfollower was aware of being unfollowed. If not, we asked whether the participant would have behaved differently if Twitter notified the unfollower. Had the participant ever broken off a relationship on other OSNs, such as Facebook, and how would one compare the different services? We included a question about company accounts and their marketing campaigns in Twitter. If the participant ever followed corporate accounts, we inquired about their participation in the campaign. The final question was to choose ten users the interviewee would *never* unfollow. Who were they? Celebrities, offline friends, or informative sources? These are indirect answers behind why people use Twitter and how they use Twitter.

Motivations behind Unfollow

The participants recalled 3.32 cases of unfollow, and some of them described multiple reasons for an unfollower. We categorize respondents’ answers in Figure 9.

Twenty out of 22 respondents reported that they unfollowed 39 people because of burst tweets. We break down these 39 into subgroups by the tweet contents. Very interestingly, respondents stated that they unfollowed 13 out of 39 because of the bursts no matter what their tweets were about. Next, respondents stated that tweets of 10 unfollowees were about uninteresting topics. The tweets of last 6 unfollowees were about the mundane details of daily life. We asked follow-up questions: why respondents did not like burst tweets? The top reasons behind the dislike of burst tweets were information overload. Respondents reported that they felt over-

	Motivation Category	Case
Burst (39)	Burst-only	13
	Uninteresting topics	10
	Mundane details of daily life	6
	Automatically generated	4
	Conversation	2
	Politics	2
	Different views	1
	Complains	1
Uninteresting topics (14)	Burst	10
	Uninteresting topics-only	4
Mundane details of daily life (11)	Burst	6
	Mundane details of daily life-only	5
Politics (8)	Different views	3
	Burst	2
	Strong opinion	2
	Politics-only	1
RT with no personality (4)	RT with no personality-only	4
Advertisement (4)	Advertisement-only	4
Automatically generated (4)	Burst	4
Different views (4)	Politics	2
	Different views-only	1
	Burst	1
Conversation (2)	Conversation-only	1
	Burst	1
Complains (2)	Complains-only	1
	Burst	1
Slang (1)	Slang-only	1
Dead account (1)	Dead account-only	1
Self-presentation (1)	Self-presentation-only	1
Language gap (1)	Language gap-only	1

Figure 9. Motivation behind unfollow for 73 unfollowees. Respondents answered multiple reasons for each unfollower

whelmed at the sheer amount of information of burst tweets. Also, too many tweets from a single followee often arrived back-to-back close in time, pushing other followees’ tweets out of the timeline. Basically, it is similar to a Denial-of-Service (DoS) attack in computer networks. Some respondents said that they stopped following celebrities, although they remained big fans of those celebrities. One stated that s/he finally unfollowed a politician because of so many tweets, although s/he started Twitter to read the politician’s tweets. Another respondent stated similar feelings:

I’m a big fan of those celebrities, and I’m interested in what they talk about. But, they wrote burst tweets

frequently, and their tweets filled my timeline. I finally unfollowed them. If they had not written too much, I would not have unfollowed them. (Respondent #8)

Four respondents counted the automatically generated log messages by third-party applications, such as Foursquare, as a cause of unfollow. Foursquare is a location-based social networking service [1]. Users with mobile phones *check in* at a venue and obtain points and badges like a game. Foursquare supports automatic transfer of user activity logs to Facebook or Twitter. The activity log typically consists of three-tuples: a short message, a location, and a Foursquare page URL. The trail of tweets by Foursquare describes a user's mobile trajectory over time. Respondents exhibited a distaste for automatically generated bursts of Foursquare logs. One respondent said:

I unfollowed a user who frequently left Foursquare log messages. If not for the Foursquare logs [or similar reports], I would not have unfollowed him/her. I am really annoyed at Foursquare logs. (Respondent #1)

Apparently many users configured their Foursquare logs to be fed to their Twitter accounts [2]. With the prevalent use of this feature, other applications also have been launching a similar feature of leaving automatic tweets, such as a new record in iPhone game, a movie or song title in a streaming service, and even a shopping list in online shopping mall. Our survey shows that these features should be used with caution. Especially, a large number of tweets in a short time could trigger unfollows.

There were two respondents who did not unfollow in the face of burst tweets. They had a different approach to browsing Twitter. They said that they had given up reading all tweets of followees. Instead, they sometimes visited Twitter, looked over tweets of a few pages only, and checked for mentions and direct messages. They thus did not care much about burst tweets. Both participants have over 400 followees.

One major outcome from the discussion of burst tweets is the need for the user interface redesign. Even when a user follow others based on similar interest, s/he would unfollow in the presence of burst tweets. Our idea is to show only one latest tweet per user and hide the rest, and upon request (for example, by placing the cursor over a 'display all' button) the rest of the tweets are displayed. This interface ensures minimizing the stress on the user from burst tweets, particularly, tweets arrived back-to-back close in time. We asked a few respondents if such an interface would help deal with burst tweets and got positive answers.

Next, 11 respondents stated that they unfollowed 14 people because their tweets were about uninteresting topics, irrespective of the quality of the tweets. Here we cast the following question: Why did you follow the person at the first place? Respondents' answers can be put into two groups. First, when Twitter sent a notification email of a new follower, they followed back if the new follower did not look like a spammer. Respondents stated that they usually fol-

lowed back a new follower. While previous work by Kwak *et al.* reports a low reciprocity of 22% in Twitter [18], our respondents surprised us: people followed back in return for being followed. The propensity for reciprocal following was developed into a solid relationship over time by shared interests, jobs, institutions, and so on. This we could call homophily in action. Second, respondents often found some tweets that appeared on their timeline via their followee's retweets and ended up following the original writer of a retweeted tweet. Retweets offered a chance to find people who wrote interestingly enough to be retweeted as social filtering. Recently, Twitter added an easy interface to retweet and now users view the retweet with the original writer's profile photo, the retweeter's name, and the total number of retweets. It helps identifying who the original writer of a retweeted tweet is.

Eight respondents replied that they unfollowed 11 people because their tweets were about the mundane details of daily life. One respondent stated:

I think Twitter is a place to exchange ideas of certain quality, while it can also be a place to express personal feelings. I, of course, am not against emotional tweets. But, some tweets to just announce personal feelings, activities, or something. They are really about mundane details. "Oh, meat is good" is completely useless to me. It is different from "Meat is good, and I wonder how it eventually got delivered to my table." Don't you agree? (Respondent #14)

Six respondents gave political tweets as a reason to unfollow 8 people. Among them, 3 unfollowees wrote tweets of differing political views. This points at two interesting interpretations. While Weng *et al.* have reported people are likely to follow those who have similar interests based on cross-sectional data [32], our respondents showed a possibility of another scenario. Users readily followed back new followers. These relationships are not always formed based on shared interests but more on curiosity and habit, and thus homophily had not yet prevailed. After a while, people got to know each others' disagreeing opinions, in this case, on politics and eventually fell apart. Homophily eventually reigned in. The other is about positioning of Twitter as a soft news medium. Baum and Jamison argue that for politically unattentive citizens soft news are more effective than traditional news [4]. Both short messages and social interaction among users could establish Twitter as a soft news medium, but Twitter users did not seem to readily listen to different voices when it comes to politics. Thus, we claim that Twitter can be an effective medium to disseminate political messages to supporters, but not to conciliate opponents due to the lack of relationships to reach them.

Finally, 4 respondents stated that they unfollowed 4 people because they frequently retweeted seemingly on random topics. Also, the same number of respondents gave advertisements, automated messages, and differing views as a reason. The rest seem to reflect individual preference of respondents.

Awareness of Being Unfollowed

After the question about the experience of unfollow, we asked whether the participant thought the unfollowee was aware of being unfollowed.

A half of respondents stated that they thought unfollowees were aware of being unfollowed. Nevertheless, they unfollowed because: (i) they did not know unfollowees in person, (ii) they got used to unfollowing although they hesitated previously, and (iii) unfollow was easy. A few respondents stated that they unfollowed even off-line friends because they used Twitter as a tool for sharing information not keeping in touch. Two of them used the metaphor of an RSS reader. The RSS reader is used for reading new articles of blogs. People readily subscribe or unsubscribe blogs without any inhibition, and they do so in Twitter.

The other half of respondents thought that unfollowees were not aware of being unfollowed. Their reasons were: (i) unfollowees had too many followers to notice, (ii) Twitter did not offer a convenient interface to track it, and (iii) they themselves did not track who unfollowed them. Our follow-up question for these respondents is “If Twitter offers a feature to notify being unfollowed, would you still unfollow others?” Respondents replied that they would not be able to unfollow at least those who they knew in person.

We note that all respondents answered that they feel free to unfollow in Twitter compared to other OSNs, such as Facebook. Most respondents did not consider the followee’s awareness seriously unless they were off-line friends or colleagues.

Company Accounts and Marketing Campaigns

Eight respondents answered that they have followed company accounts for marketing campaigns, and five of them kept following. The expectation for continual benefit was one of the reasons to keep following. One respondent stated:

I love those events. I have followed a few company accounts. [After marketing campaigns end,] I do not unfollow them because they may do another marketing campaign in the near future. I unfollow company accounts only if they leave too many tweets. (Respondent #2)

Another reason was appreciation for an opportunity to win prizes. The rest of the three stated that they found no need to unfollow company accounts, because those accounts generated only a few tweets. All 5 respondents answered that they did not mind reading a few advertisement tweets occasionally, but all those users assured us that they would certainly unfollow company accounts if they leave burst tweets.

Twelve respondents answered that they have never followed company accounts conducting marketing campaigns. Their reasons were various. Some did not want to read even one advertisement tweet. They carefully chose whom to follow and did not want uninteresting tweets. Other respondents showed a strong sense of responsibility toward retweets. The

most common form of a marketing campaign in Twitter is retweeting advertisement tweets, and respondents did not want to deliver an advertisement tweet to their own followers. These users retweeted only those tweets they deemed important enough to be disseminated broadly. Some of respondents stated that they did not like a lottery system for win prizes; they only participated in marketing campaigns where all participants received a gift.

Who Would You Unfollow Last

Finally we questioned participants on who would be the last ten people to unfollow. While Kwak *et al.* have highlighted that Twitter has the characteristics of news media rather than social networks, most respondents chose intimate friends as the last 10 people to unfollow. They wanted to read tweets of strong ties rather than those of informative sources. Even though the Korean Twitter network still has a low reciprocity hovering slightly over 50%, at its core the strong ties are what keeps the people connected. This evidence connects to the study of Huberman *et al.* [13]; they have revealed that each user has small number of friends, those who exchange at least two tweets, compared to declared followees. Some of respondents chose those who were role models to them, e.g. CEOs, politicians, and professors. They enjoy reading tweets capturing their role models’ view. Only on Twitter they could access such information of their role models. Respondents stated that they found role models’ tweets inspirational.

CONCLUSION

We have explored the dynamics of the unfollow behavior in Twitter. We collected daily snapshots of the online relationships of 1.2 million Korean-speaking users for 51 days and their all tweets. We observed the prevalent unfollow behavior in Twitter. We found that the major factors, including the reciprocity of the relationships, the period of the relationships, the followee’s informativeness, and the overlap of the relationships, are crucial for the decision to unfollow. We conducted interview with 22 users to determine their motivation behind the unfollow behavior. The survey participants unfollowed those who left many tweets within a short time, created tweets about uninteresting topics, or tweeted about the mundane details of their lives.

ACKNOWLEDGEMENT

We are grateful to Hyunyoung Song, Alice Oh, Youn-kyung Lim, Sangjin Han, Seungyeop Han, Keon Jang, DK Lee, and Changhyun Lee for thoughtful feedback. We would like to thank all respondents for participating in survey. We also thank anonymous reviewers for their constructive and generous suggestions. This work was supported by the IT R&D program of MKE/KEIT [KI001878, “CASFI : High-Precision Measurement and Analysis Research”].

REFERENCES

1. Foursquare. <http://foursquare.com>.
2. Foursquare logs in Twitter. <http://twitter.com/#search?q=4sq.com>.

3. A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.
4. M. Baum and A. Jamison. The Oprah effect: How soft news helps inattentive citizens vote consistently. *The Journal of Politics*, 68(4):946–959, 2006.
5. H. R. Bernard, P. Killworth, D. Kronenfeld, and L. Sailer. The problem of informant accuracy: The validity of retrospective data. *Ann. Rev. Anthropol.*, 13:495–517, 1984.
6. R. B. Cairns and B. D. Cairns. *Lifelines and Risks: Pathways of Youth in Our Time*. Cambridge, 1995.
7. H. Chun, H. Kwak, Y.-H. Eom, Y.-Y. Ahn, S. Moon, and H. Jeong. Comparison of online social relations in volume vs. interaction: A case study of Cyworld. In *Proc. IMC*, pages 57–70, 2008.
8. Economist.com. Primates on Facebook: Even online, the neocortex is the limit, <http://goo.gl/IxEsq>. *The Economist*, Feb 2009.
9. N. B. Ellison, C. Steinfield, and C. Lampe. The benefits of Facebook friends: social capital and college students use of online social network sites. *Journal of Computer-Mediated Communication*, 12(4):1143–1168, 2007.
10. Y.-H. Eom, C. Jeon, H. Jeong, and B. Kahng. Evolution of weighted scale-free networks in empirical data. *Phys. Rev. E*, 77(5):056105, 2008.
11. S. A. Golder and S. Yardi. Structural predictors of tie formation in Twitter: Transitivity and mutuality. *Proc. SOcialCom*, pages 88–95, 2010.
12. M. T. Hallinan and R. A. Williams. The stability of students’ interracial friendships. *American Sociological Review*, 52(5):653–664, 1987.
13. B. Huberman, D. M. Romero, and F. Wu. Social networks that matter: Twitter under the microscope. *First Monday*, 14(1-5), 2009.
14. A. Java, X. Song, T. Finin, and B. Tseng. Why we twitter: understanding microblogging usage and communities. In *Proc. KDD/SNA*, pages 56–65, 2007.
15. D. B. Kandel. Homophily, selection, and socialization in adolescent friendships. *American Journal of Sociology*, 84(2):427–436, 1978.
16. G. Kossinets and D. J. Watts. Origins of homophily in an evolving social network. *American Journal of Sociology*, 115(2):405–450, 2009.
17. R. Kumar, J. Novak, and A. Tomkins. Structure and evolution of online social networks. In *Proc. KDD*, pages 611–617, 2006.
18. H. Kwak, C. Lee, H. Park, and S. Moon. What is Twitter, a social network or a news media? In *Proc. WWW*, pages 591–600, 2010.
19. C. Lampe, N. B. Ellison, and C. Steinfield. Changes in use and perception of Facebook. In *Proc. CSCW*, pages 721–730, 2008.
20. K. Lerman and R. Ghosh. Information contagion: an empirical study of the spread of news on Digg and Twitter social networks. In *Proc. ICWSM*, 2010.
21. J. Leskovec, J. Kleinberg, and C. Faloutsos. Graphs over time: densification laws, shrinking diameters and possible explanations. In *Proc. SIGKDD*, pages 177–187, 2005.
22. D. Liben-Nowell and J. Kleinberg. The link-prediction problem for social networks. *Journal of the American Society for Information Science and Technology*, 58(7):1019–1031, 2007.
23. D. A. Menchik and X. Tian. Putting Social Context into Text: The Semiotics of E-mail Interaction. *American Journal of Sociology*, 114(2):332–370, 2008.
24. M. B. Miles and A. M. Huberman. *Qualitative Data Analysis: An Expanded Sourcebook*. SAGE Publications, 1994.
25. G. Mollenhorst. *Networks in contexts: How meeting opportunities affect personal relationships*. PhD thesis, University Utrecht, 2009.
26. J.-P. Onnela, J. Saramäki, J. Hyvönen, G. Szabó, D. Lazer, K. Kaski, J. Kertész, and A.-L. Barabási. Structure and tie strengths in mobile communication networks. *PNAS*, 104(18):7332–7336, 2007.
27. C. Steinfield, N. B. Ellison, and C. Lampe. Social capital, self-esteem, and use of online social network sites: A longitudinal analysis. *Journal of Applied Developmental Psychology*, 29(6):434 – 445, 2008.
28. S. T. Tong, B. Van Der Heide, L. Langwell, and J. B. Walther. Too much of a good thing? the relationship between number of friends and interpersonal impressions on facebook. *Journal of Computer-Mediated Communication*, 13(3):531–549, 2008.
29. Z. Tufekci. Who acquires friends through social media and why? “rich get richer” versus “seek and ye shall find”. In *Proc. ICWSM*, 2010.
30. N. B. Tuma and M. T. Hallinan. The effects of sex, race, and achievement on schoolchildren’s friendships. *Social Forces*, 57(4):1265–1285, 1979.
31. S. Wasserman and K. Faust. *Social Network Analysis: Methods and Applications*. Cambridge University Press, 1994.
32. J. Weng, E.-P. Lim, J. Jiang, and Q. He. TwitterRank: finding topic-sensitive influential twitterers. In *Proc. WSDM*, pages 261–270, 2010.
33. D. Zhao and M. B. Rosson. How and why people Twitter: the role that micro-blogging plays in informal communication at work. In *GROUP*, pages 243–252, 2009.