Do I Need To Follow You? Examining the Utility of The Pinterest Follow Mechanism

Bluma Gelley*
Computer Science and Engineering
New York University
Brooklyn, NY, USA
bgelley@nyu.edu

Ajita John
Avaya Labs
Basking Ridge, NJ, USA
ajita@avaya.com

ABSTRACT
Pinterest is a Social Network Site (SNS) centered around the curation and sharing of visual content. The site encourages users to form ties with (follow) other users based on mutual interests, and use these ties to discover and share content. In this work, we examine the efficacy and relevance of the Pinterest follow mechanism in driving content discovery and curation. We collect a sample of user activity and find that the vast majority (88%) of the unique users who interact with an average user's content are non-followers. Conversely, only 12.3% of a user's followers interact with any of their pins. Users who discover and repost content from outside their follow network also do not subsequently follow the contributors of that content. Our results strongly suggest that following is neither heavily utilized nor strongly effective for driving content discovery and sharing on Pinterest.

Author Keywords
Pinterest; Online Social Networks; Following; Social Browsing; Activity Graph; Tie formation

ACM Classification Keywords
H.3.5. Information Systems: Online Information Services—*Web-based Services*; J.4.3. Social and behavioral sciences

General Terms
Human Factors; Measurement

INTRODUCTION
The follow mechanism — the public articulation of links between users[1] [6] — is an essential part of a Social Network Site (SNS) [7]. According to boyd and Ellison [21], this mechanism is "the key differentiating feature of SNSs." Following is what connects users to the social graph maintained on the site and allows them access to the many benefits it provides [52, 65]. Follow links are the means by which users influence others [62, 2] and ideas and advertisements are spread through the network [65, 12, 13]. A user's follow network can even be used to target ads for the products and services that they may be likely to purchase [31]. As such, the follow mechanism is a crucial part of any social network site.

The act of following (or friending) another user can be understood as an example of the sociological concept of tie formation in social networks [27]. Various mechanisms for tie formation have been well-studied in the sociology and psychology literature, including homophily, friendship based on similarity or mutual interests [45]; triadic closure, connecting with friends of friends [57]; and preferential attachment, or the tendency of people to connect with others who already have many relationships [3]. People form ties on social network sites as well for similar reasons [57, 54, 8]. In the early days of the Internet, communities were formed by people who shared common interests; one example of this is Usenet newsgroups. Tie formation in those communities resulted from ***interest homophily*** [4] — connecting with others because of common interests. Those early communities, however, never gained the mainstream acceptance that later social network sites such as Facebook have [21]. These sites "signaled a shift from interest-driven to friendship-driven spaces" [21], generally depending on mechanisms other than interest homophily to build their networks. This shift is substantiated by the work of Bisgin and colleagues [4], who found that shared interests had only a weak influence on tie formation on several mainstream SNSs of varying types. On most modern social network sites, the majority of ties that are 'formed' are actually ones that already existed off the site. A Nielsen survey of Internet users showed that the top reason given by users for friending someone on Facebook was "knowing [them] in real life" [35]. The 2013 Surveying the Digital Future survey also found that most users had only a small number of friends whom they had met online [16]. This has been confirmed repeatedly for a number of social network sites: the majority of a user's connections are usually 'real life' friends and acquaintances who were simply transferred to the site as 'friends' [22, 39, 41, 32, 14, 23]. This transfer is often encouraged by the site: Facebook users are encouraged to use the site's search feature to find people they already know and add them as friends. LinkedIn takes this idea even further by

*A significant portion of this work was done while the author was at Avaya Labs.

[1]In this paper, for purposes of simplicity, we use the word 'follow' to refer to both unidirectional ('follow') and bidirectional ('friend') links on a social network site.

CSCW 2015, March 14–18, 2015, Vancouver, BC, Canada.

ACM 978-1-4503-2922-4/15/03 ...$15.00.
http://dx.doi.org/10.1145/2675133.2675209

allowing users to connect only with people they already know. Users attempting to add another user to their network are required to specify exactly how they are previously connected to the person; otherwise, they must provide the person's email address or ask to be introduced by a mutual acquaintance.

Recently, however, an SNS with a tie-formation strategy based on interest homophily has gained prominence: Pinterest. Pinterest is a social network site centered around the curation and sharing of visual content. Since its inception in 2010, it has grown extremely rapidly, reaching 10 million monthly unique visitors faster than any SNS ever[17] and boasting 70 million users by June 2013 [55]. A 2013 Pew Internet Survey [20] found that 21% of all Internet users in the US use Pinterest. By early 2014, Pinterest was driving 7.1% of all referrals to purchasing sites, second only to Facebook [19]. Pinterest has attracted much attention from marketers, due to its 'aspirational' nature [49, 47], with users using the site to find and share products and services that they would like to buy. The site uses the metaphor of virtual 'pinboards' on which images and other media (known as ***pins***) are 'pinned'. An example of a ***board*** (a collection of pins on a specific topic) is found in Figure 1.

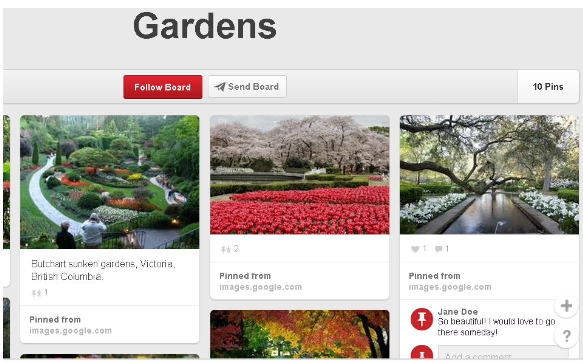

Figure 1: An example of a Pinterest board in the 'Gardening' category (see next section).

Hall and Zarro [29] use the term ***social curation*** to describe the activity on Pinterest. Users collect pins, sometimes from the boards of other Pinterest users, sometimes from elsewhere on the Internet, and arrange them into topical collections. This process of selection and categorization adds value to the individual pieces of the collection and is referred to as ***curation*** [58]. The social element is introduced through the publicly viewable nature of these collections and the ability of users to "view, favorite...copy, and comment" [29] each others' content. Pinterest's focus on the *curation* of others' content is what separates it from other content-focused SNSs like Flickr and YouTube, which center around the creation and sharing of user-created media [61, 46].

Like other social network sites, Pinterest is built around the concept of a social graph, allowing users — known as ***pinners*** — to ***follow*** other users and be shown the content they post. Pinterest uses the ***social awareness stream*** [51] mechanism popularized by sites like Facebook and Twitter, called the ***home feed*** on Pinterest, which shows users all new content posted by the users they follow as it is posted. This method of finding relevant content by browsing through the content posted by one's friends has been labeled ***social browsing*** by Lerman and Jones [42]. Follow links on Pinterest are unidirectional — users can follow other users who do not follow them.

Unlike the relationship-centered SNSs described earlier, Pinterest relies on interest homophily to build its social graph. The site is centered on content, with all activity revolving around pins. Users are encouraged to follow others whose content they find interesting, rather than people they know offline. In fact, the recommended follow mechanism on Pinterest is following just specific boards, rather than whole users.[2] In the words of a Pinterest user quoted by Zarro, Hall, and Forte [68], Pinterest is a "community of people who don't know each other." In this way, Pinterest is attempting to create a new social graph, based not on pre-existing offline connections, but on mutual interests. While Pinterest, like Facebook, does provide a mechanism to find people on the site whom a user already knows, Pinterest carefully limits this to "friend[s] who share your interests."[3] Facebook, on the other hand, suggests that users "send a friend request to any of your friends that already have a Facebook account" and even invite to the site those who don't.[4] Pinterest further cements its emphasis on content as a basis for following by asking new users what topics they are interested in and presenting them with a selection of relevant boards to follow, rather than users.

On the surface, the interest-homophily tie-formation model seems to be highly successful on Pinterest. We found that the median number of followers per user is 106, more than both Facebook and Twitter at a later stage of their growth [60, 2]. In addition, Chang et al. [11] report that there is indeed a high degree of interest homophily in the Pinterest follow network. However, a social graph has little value on its own; its success can only be measured by how well it advances the goals of the site and its users — in this case, discovering, curating, and sharing visual content [11]. To the best of our knowledge, no one has yet examined the effectiveness of the Pinterest follow mechanism in achieving these goals.

In this paper, we examine following on Pinterest, exploring some intriguing questions. How effective is the Pinterest follow model at fostering content discovery, the spread of ideas, and interaction based on common interests? Following is essential to Pinterest's success as a social network site, but is it essential to the average Pinterest user's experience with the site? Can a follow mechanism based on interest homophily succeed on a large-scale, mainstream SNS?

To answer these questions, we turn to the ***activity graph***, the hidden network formed by interactions between users. A large body of work has shown that the network formed by linking users with their followers — the ***follow graph*** — contains incomplete information about users and their relationships. Studies of Facebook [64], Twitter [33], and Cyworld

[2]Though in practice most users choose to follow whole users; see the next section.

[3]http://help.pinterest.com/en/articles/find-your-friends

[4]https://www.facebook.com/help/336320879782850/

[15] all found that the activity graph is very different from the follow graph. Specifically, users tend to interact with only a small subset of their friends and followers [44, 33]; thus, relationship strength, and in turn, strength of influence, are difficult to predict using the follow graph alone [10, 59, 56]. On Pinterest, activity takes the form of liking, commenting on, and repinning other users' pins. These actions are the means through which content is curated and shared. Repinning is a curation action — reposting desired content in the user's own space. Liking can be a curation action as well, since users often like pins to save them for later, as described in the next section. Comments have less of a curation dimension and are therefore rare; all three are indications that other users have seen and reacted to the content. Activity can thus be viewed as traces that users leave behind when they view content; it can be used to track user viewing patterns that would otherwise be invisible. We therefore sample the implicit activity graph on Pinterest and compare it to the follow graph to determine whether the follow mechanism - through the medium of social browsing - is successfully promoting content discovery, curation, and sharing. To guide our examination, we formulate the following research questions.

RQ1: What is the relationship between followers and activity?

If following is indeed successful at fostering activity, we would expect to find a strong relationship between a user's number of followers and the amount of activity done to their pins. Does such a correlation indeed exist? Do followers and activity follow similar patterns in their relationships with other metrics, such as a user's number of pins?

RQ2: How heavily do followers interact with their followees' content?

If the follow mechanism is heavily utilized for content discovery, we would also expect most of a user's followers to engage with at least some of their content. We therefore ask: are most followers active on their followees' boards? If so, to what extent?

RQ3: Do users find content from sources other than their followees?

Do users primarily utilize their home feeds when finding interesting content, and therefore only see (and curate) content from the users they follow? Or do they make use of other sources for pins as well?

RQ4: Do users follow the 'out-of-network' users or boards whose content they interact with?

Pinterest provides several mechanisms outside the home feed to help users discover content from pinners they do not follow, including search and featured content pages. The site encourages users to follow the boards or pinners where they found the content. This will allow them to see more, potentially similar, content from the user or board in the future. Do users avail themselves of this option?

The answers to these questions may shed some light on the importance of the follow mechanism for content discovery and curation.

We present the following results:

- Though following is supposed to be based on mutual interests, only a small percentage of a user's followers interact with the user's content or add any of it to their own collections. On average, only 12.3% of a user's followers like, comment, or repin any of their pins.
- A large majority (>70%) of the activity on a user's posts comes from users who are not their followers. Those followers who do interact, however, perform nearly 2.5 times more actions, on average, than non-followers. This suggests that users find content to curate in locations other than their home feeds or the boards of those they follow.
- Users who collect content from the Popular or Category pages do not generally proceed to follow the pinners of that content or their boards, though Pinterest encourages them to do so; an average Popular pin receives 673 likes, comments, and repins, but the user who posted it gains just one new follower for every 450 of these interactions; more than half gain no followers at all.
- Taken together, our results suggest that following is neither heavily utilized for nor strongly effective at promoting content discovery, curation, and sharing on Pinterest.In this way, Pinterest differs significantly from Flickr, another SNS focused on visual content, where following and the resultant social browsing has been shown to be important for content discovery [42].

The remainder of this paper is structured as follows. In the Background section, we provide an overview of Pinterest and define the terms we use in the rest of the paper. We then discuss related work. We devote the next section to describing our sampling and data collection methods, along with the resulting dataset. In the Data Analysis section, we report the results obtained from our analysis of the dataset; we then discuss their implications in a section entitled Discussion. Finally, we conclude and propose some ideas for future work.

BACKGROUND

Overview of Pinterest

Pinterest describes itself as "a place to discover ideas for all your projects and interests, hand-picked by people like you."[5] It's billed as a 'virtual pinboard' service, where users can easily 'pin' digital content they find interesting or useful and share it with others. The central entity on Pinterest is the ***pin***. A pin is an image or video, often accompanied by a caption. Pins can be uploaded by the user, but the vast majority are reposted from somewhere else on the Web [24]; these link back to the original source when clicked on. Users, or ***pinners***, ***pin*** content onto their ***boards*** — pages, usually organized around a specific theme, where pins are laid out in an informal style reminiscent of a physical pinboard. Pinterest's trademark layout is designed for maximum visual appeal: pins are displayed in neat rectangles of varying heights in a grid pattern that continuously loads new content as the viewer scrolls down. Clicking on a pin opens it on a separate

[5] https://about.pinterest.com/en

page with more detail. See Figure 1 in the Introduction for an example of a board.

User profiles on Pinterest are relatively basic: users can add a profile picture, a brief description, and their location, as well as links to Facebook, Twitter, and/or a personal website. Also displayed are the user's number of boards and pins and the number of pins they have liked, as well as the number of other users they follow (called 'following') and the number of followers they have. The rest of the profile page is devoted to the user's boards. Like pins, the board 'teasers' on a user's profile page are laid out as small rectangles in a grid, and display several images from the board. Every board on Pinterest belongs to a topical category (or 'Other'), and the pins on each board are supposed to fit that topic.

Pinterest users can connect with other users by ***following*** them. New pins from followed boards or users show up in a follower's ***home feed***, similar to a Facebook news feed, which they see when they open Pinterest. Follow edges in Pinterest are directed; pinners can follow other users who do not follow them back, and following someone does not require their permission. Since pinners often have boards on many different topics, users are encouraged to follow only those boards which interest them. However, Ottoni et al. found that in practice, users overwhelmingly (~90%) choose to follow other pinners entirely rather than single boards [53]. As well, despite the site's emphasis on following boards over users, board followers and followees are counted in the user's general follower and followee counts on his/her profile page. Users can also create and join group boards, some of which have thousands of pinners posting content to them. Each user is also allowed up to 3 secret boards, which are only visible to the owner(s).

There are three main types of interaction on Pinterest: likes, comments, and repins. According to the Pinterest help page, "Like a pin when you want to say Hey you! Neat idea!" When a user likes a pin, it is also saved to the 'Likes' page in her profile, so pinners will often like pins that they want to be easily able to find later. Users can also comment on a pin; comments are displayed beneath the pin on the board. Finally, users can repin a pin to one of their own boards. Since Pinterest's goal is primarily *curation* of content [25], rather than creation, repins are treated no differently than ordinary pins: they are displayed on the user's board and in her followers' home feeds without even an indication that they were originally pinned by someone else. We use the term post to include both pinning and repinning. Any logged-in user can like, comment on, or repin any other user's pins, without needing to follow the user or board. We visualize the Pinterest data model and the interactions between entities in Figure 2.

Pinterest provides several methods for content discovery besides the home feed. A continuously changing sample of pins from boards in each category are reposted on the Category pages, accessible from a drop-down menu on the site header. There is also a 'Popular' page, where popular pins from around the site are displayed, and a search box for finding pins, boards, or users matching a search term.

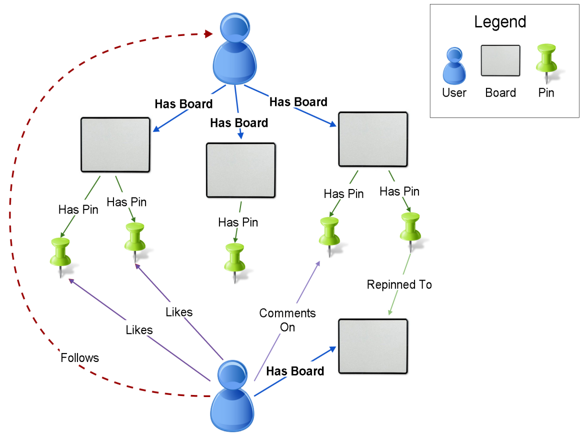

Figure 2: The Pinterest data model. Note that it is not necessary to follow a user in order to repin, like, or comment on their pins.

Definitions

In this section, we define some of the terms we use.

Follow graph: a graph representation of the users in a social network where vertices represent users and the edges between users are formed by users following (directed) or friending (undirected) each other

Activity graph: a graph representation of the users in a social network where the directed edges between users represent interactions between them. Two users A and B are connected in the activity graph if there was an interaction of some sort between them. ***Activity link***: an edge in the activity graph, formed by an interaction between the users at its ends.

Follow link: an edge in the follow graph, formed by a friend/follower relationship.

Pinterest activity link: a like, comment, or repin *received by* a user. Though the interaction is technically with the pin itself, it creates a link to the pin's owner. In this work, we calculate per-user statistics based on *incoming* activity — actions done *to* a user's content, rather than *by* the user.

Post: A pin or repin. These are not distinguished in Pinterest, but we use this term here for clarity.

RELATED WORK

Pinterest

Pinterest is a fairly new site, and its lack of an API has created an additional barrier to its study. A few analyses of Pinterest, however, have been published very recently. Gilbert et al. [25] attempt to determine what drives user behavior on Pinterest by calculating the contribution that various factors (such as the gender and nationality of the original pinner) have to the likelihood of a pin's being repinned and the number of followers a user attracts. They also compare the language used by the same users on Pinterest and on Twitter and determine that there are significant differences. Feng et al. [24] study user behavior on Pinterest, but they confine their analysis to the static follow graph; they also study the content and categories of pins, particularly popular ones. Mittal et al. analyze various aspects of Pinterest, including some user characteristics, the distribution of user locations, and pin sources. They also address privacy and copyright issues and find many instances of personal data leakage and copyright violations on Pinterest. Finally, they find that they can predict gender of Pinterest users with high accuracy [48]. Gender differences in Pinterest are a popular topic of study; Ottoni et al. [53] quantify the differences in Pinterest behavior between male and female users. Chang and colleagues [11] also study gender, specifically, the types of content favored by, and degree of specialization of, the two genders. They also report that homophily — here defined as similarity in interests — has a large influence on repinning, but a smaller one on following. Kamath et al. [36] build a model to automatically recommend boards that users might like. Zarro et al. [68] studied Pinterest from a qualitative perspective through user interviews. They found that users see Pinterest as a content provider rather than a social network. This was confirmed by the findings of Han et al. [30], who showed that repinning is mostly driven by the pin's content rather than user characteristics.

Activity Graph

The concept of the implicit activity network in an online social network and the fact that it differs from the explicit follow network was first proposed by [15], who analyzed the topographical characteristics of both the follow and activity networks of Cyworld, a Korean SNS. They found that the one-way interaction network had a similar topology to the follow network, but the reciprocal 'friends' network was quite different, more similar to known topologies of offline social networks than to the usual characteristics of online social networks. Ahn et al. [1] had previously made a similar observation about the testimonial network on Cyworld, but did not extend their results to the concept of the activity graph in general. Wilson and colleagues [64] performed a very similar analysis on Facebook, referring to the implicit network as the interaction graph.They found significant differences between the follow graph and the interaction graph, once again finding that the interaction graph displays the small-world properties typical of online social network graphs to a lesser extent than the follow graph does. Viswanath, et. al. [63] studied the evolution of activity links over time and discovered that Facebook activity links change over time, but many of the graph-theoretic properties of the activity graph did not. The Twitter interaction graph was studied by Huberman et al. [33], who found that most users interact closely with only a small subset of their followers. This disparity was confirmed in the case of Facebook by the Facebook Data Science team, who, with access to all of Facebook's user data, showed that the number of active reciprocal relationships per user was much smaller than the user's friend count [44].

Tie Formation and Homophily in Social Networks

A large body of research exists in sociology and psychology on tie formation and homophily, too numerous to cite here. We refer readers to McPherson et al.'s literature review on homophily [45], and Donath and boyd's discussion of some previous work on tie formation [18]. Here we briefly list research on homophily and tie formation in online social networks. Lauw et al. found that users on LiveJournal with similar interests were more likely to be friends [40], but the much larger study of Bisgin et al. found no evidence of interest homophily on several social network sites [4]. Macskassy and Michelson found that models that take homophily into account better fit retweet behavior on Twitter than those that don't [43]. Kivran-Swaine et al. [37], Quercia et al. [54], and Kwak et al. [38] studied the breaking of ties on social network sites. Romero and Kleinberg looked at the process of directed closure in social networks through the lens of follow links on Twitter [57]. There is also a large amount of work on modelling ties [67] and predicting link formation in SNSs, including [26, 34, 27].

DATA COLLECTION & DESCRIPTION

As discussed in the introduction, we concentrated on the activity graph rather than the follow graph. Instead of crawling follower edges, we followed activity links between users (though we did collect follow links as well for crawled users). As is common in social network analysis, analyzing the entire Pinterest graph was impractical; we therefore did our analyses on a sample of the network. Since unique ids on Pinterest are text-based, random sampling was difficult. In addition, Pinterest lacks an API, and employs various design techniques that make collecting all data for each user prohibitively expensive. We therefore used a modified Breadth-First Search (BFS) (that is, Snowball Sampling [28]) on the full graph to collect a sample of nodes and edges. We began the crawl from several randomly chosen seeds and moved outward by selecting a random subset of edge clusters and crawling all edges in each cluster. This is accomplished by randomly choosing 5 boards of each crawled user, and then crawling all activity links on each board. Likewise, we limited the number of pins collected per board to the first 300, since 90% of boards have 300 or fewer pins. On a similar SNS graph, Bonneau et al. [5] were able to estimate many graph properties using, as we do, a random sample of k edges from each node, even with very small k. Due to the extreme difficulty and expense of crawling large follower lists, we limited our analysis of activity from followers (see the end of the Data Analysis section) to users with 10,000 followers or fewer; this includes 99.96% of the users in our dataset.

We utilized a multithreaded crawler architecture for data collection. We ran the crawler for 5 weeks in December 2013 and January 2014 and collected the data described in Table 1. The first row consists of users whose incoming activity we sampled as described previously. For those users who liked or commented on one of the crawled pins, or followed one of the crawled users, we collected their real names, number of pins and number of followers (here denoted 'partial data'); 'users touched' includes the two categories above it, as well as users who repinned one of the pins in our dataset, for whom we have only their usernames. Boards are similar — we have full metadata, plus up to 350 pins, for the crawled boards, partial data (URL and number of pins) for an additional 2 million boards, and just the URL for the final ~2.9 million.

Dataset Details	
Crawled Users (with boards and pins)	31,644
Users (partial data)	4.5 million
Total Users touched	5.4 million
Crawled Boards	163,300
Boards (partial data)	2 million
Total Boards Touched	5.1 million
Total Pins Crawled	14 million
Total Repins	7 million
Total Likes	1.56 million
Total Comments	47,557

Table 1: Data Description

	Mean	Med.	Mode	Stdev	Max
Boards per User	30.1	19	12	41	2,310*
Pins per Board	138	39	1	563	100,228**
Followers per User	604	64	0	26,639	4,283,442

*All are group boards.
**This is the largest number of pins on a board with a single pinner. Group boards can have millions of pins.

Table 2: Statistics for the dataset. All mins are 0.

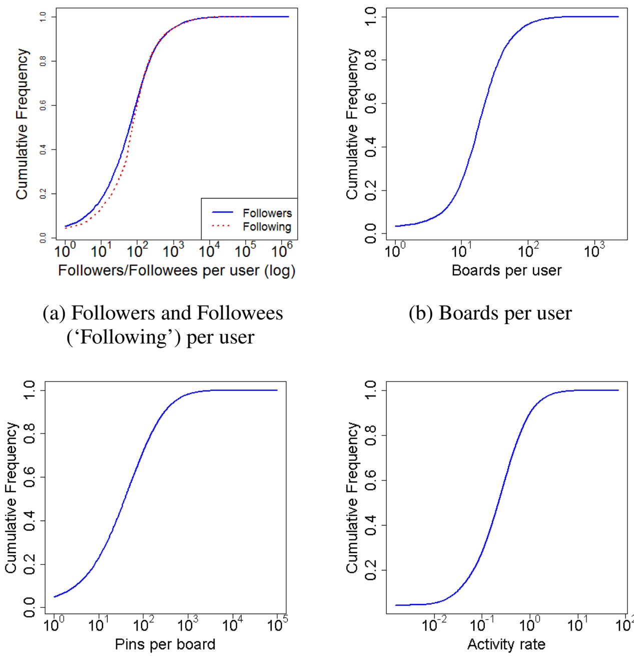

(a) Followers and Followees ('Following') per user

(b) Boards per user

(c) Pins per board

(d) Activity Rate per user

Figure 3: Cumulative frequency distributions of followers and followees (3a), boards (3b), and activity rate (3d) per user, and pins per board (3c), in log scale.

Descriptive Statistics

We first present some descriptive statistics about our dataset, including distribution of the three types of activity across users, boards, and pins. Table 2 contains basic statistics for numbers of boards, pins, and followers/following; corresponding distributions are shown in Figure 3.

Table 3 summarizes the activity links contained in our dataset. Figure 4 shows the distributions of likes, comments, and repins across users, boards, and pins. While overall, there are 58 interactions (likes, comments, or repins) for every 100 pins, only 17.7% of pins have even a single interaction, due to the skew in the distribution of activity per pin. Repins are by far the most common type of activity, about 4 times more common than likes and nearly 150 times more than comments. They make up 81% of total activity recorded and, on average, 79% of activity on each board. Comments are extremely rare; just .3% of pins have even one comment. By contrast, 15.2% of the pins in our dataset were repinned at least once. We attribute this to the nature of Pinterest itself, where the primary goal is discovering and curating content (i.e. pins), with social interaction coming in a distant second. Comments by nature involve far more social interaction than repins; repinning just means that the repinner wants the content for herself, while commenting is usually a form of communication with others. In this aspect, Pinterest is different from Facebook, where liking is many times more common than resharing (even commenting is more common than resharing), and similar to Twitter, where resharing is more common than liking [50]. All of the distributions are very heavy-tailed, as is clear from the plots. The distribution of activity per pin is the most skewed; there are a few pins that seem to be more desirable or interesting, while others appear to be ignored. All of these distributions were so skewed that they had to be plotted on a log-log scale to be clearly visible.

	Mean	Med.	Stdev	Max
Activity per user	265.4	67	951.7	78,336
Activity per board	51.5	5	262.6	32,640
Repins per pin	0.47	0	4.4	1424
Likes per pin	0.11	0	1.5	755
Comments per pin	0.003	0	0.1	126

Table 3: Statistics for activity distribution. All mins and modes are 0.

Activity Rate

In this work, we use activity links as edges in the activity graph. When comparing activity between users, however, or even when comparing to the follow graph, the raw activity count is not a useful metric. Different users have different numbers of pins, so a large amount of activity may just be a function of having many pins. Activity (specifically reposting) per post is used as a measure of influence on Twitter by [59], and on Pinterest by [25]. We therefore use activity rate throughout this paper instead of raw activity

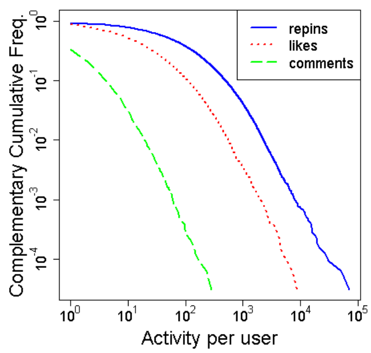

(a) Activity per user

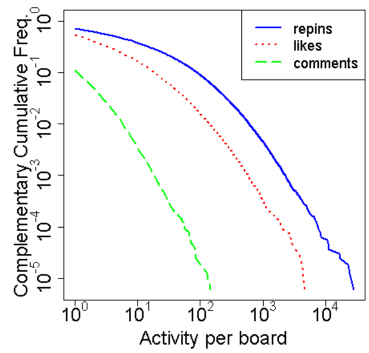

(b) Activity per board

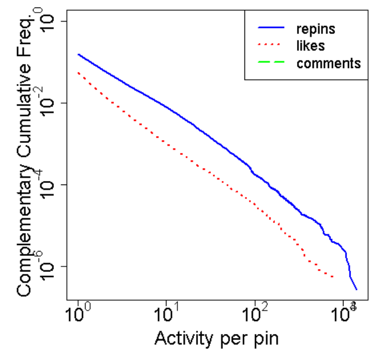

(c) Activity per pin. Comments were so rare that they did not show up in the plot, even in log-log scale.

Figure 4: Complementary cumulative frequency distributions of activity per: user (4a), board (4b), and pin (4c). The distributions are plotted on a log-log scale so that they can be shown on the same plot; despite this, the number of comments per pin is so small that the distribution line disappears off the side of even a log-log plot.

counts. As would be expected from the activity distributions shown above, activity rates vary widely. Some users have large amounts of activity per pin, while others have next to none: the minimum activity rate in our sample is 0, while the maximum belongs to the Pinterest account of the popular children's retailer Carter's — the brand averages 105 likes, repins, or comments per pin. The distribution of activity rate for the users in our sample is shown in Figure 3d.

DATA ANALYSIS

Followers and Activity — RQ1

We now begin to address our research questions, beginning with RQ1, the relationship between number of followers and activity rate. A strong relationship between followers and activity would suggest that users are engaged with the content of their followees. Figure 5a shows the number of followers per user plotted against the user's activity rate, both on a log scale so differences are clearly visible. Each hexagon represents multiple points, as shown on the key to the right of the plot, with lighter areas representing more data points. Spearman's ρ correlation coefficient between the number of followers and the activity rate is 0.55, and the R^2 goodness-of-fit is 0.26.[6] The above correlation includes users with many hundreds of thousands or millions of followers who bear little resemblance to 'ordinary' pinners; we therefore exclude the top 5% of users by follower count (> 924 followers), as well as users with fewer than 5 followers. The scatterplot of followers and activity rate for these 'average' users is shown in Figure 5b. Here, the correlation is $\rho = 0.44$, with an R^2 of 0.19. These correlations are moderately high; we therefore continue to further analysis.

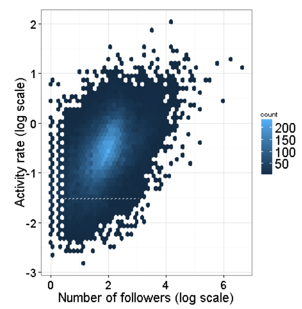

(a) Followers and activity rate for all users.

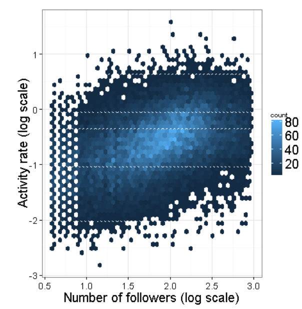

(b) Followers and activity rate for users with 5–924 followers.

Figure 5: Followers and Activity for: 5a) all users; 5b) the middle 85% of users by follower count. Each hexagon contains the amount of data points indicated by its color; lighter areas have more data points.

To shed more light on the relationship between following and activity rate, we compare each one with a third measure: pin count. If following and activity rate parallel each other closely, we would expect to see similar relationships between each of them and pin count. Figure 6a plots number of pins against follower count for each user in our sample. Spearman's ρ here is 0.78, and R^2 is 0.3. Figure 6b shows the activity rate plotted against the number of pins, using the same methodology as Figure 6a. Here, Spearman's ρ is 0.32, and R^2 is 0.06.

The relationship between pins and activity rate is much weaker than between pins and followers. This can be seen on the scatterplot; both the entire mass of data points and the lighter (denser) area are nearly circular, denoting a weak relationship between the two variables. This difference casts some doubt on the existence of a strong relationship between

Pins, Followers, and Activity Links

[6] All correlations and R^2 values reported in this and the next subsections are statistically significant, $p < .0001$.

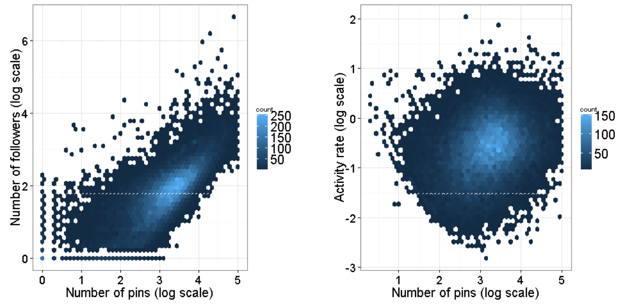

(a) Pins and followers.

(b) Pins and activity rate.

Figure 6: Number of pins (x axis) plotted against edges in the follow graph (6a) and the activity graph (6b). (All axes are log scale.) Each hexagon contains the amount of data points indicated by its color; lighter areas have more data points.

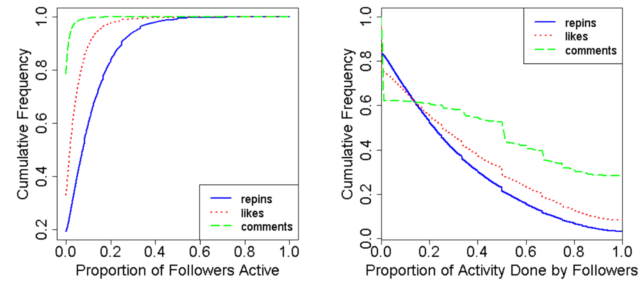

(a) Proportion of a user's followers who engaged in activity on the user's boards, by type of activity.

(b) Proportion of the activity on each user's boards done by their followers.

Figure 7: Activity by followers.

followers and activity. There does seem to be a relationship between pin count and follower count. It may be that people with more followers are more motivated to pin, or conversely, that users with more pins garner more followers. (Or there may be other factors causing both.) As for pins and activity rate, a correlation of 0.32 on SNS data is strong enough to confirm Wu et al.'s [66] finding of a positive feedback loop between feedback on posts and posting frequency. However, the relationship is clearly weaker than the correlation between pins and followers. Are pinners more motivated by follower count as a positive feedback measure than by activity? Or, perhaps, do users like following pinners with many pins but are less interested in acquiring (i.e. repinning) their content? These questions provide an intriguing base for further study.

Similar results have been reported for Twitter: Suh et al. [59] found a linear relationship between retweets and followers, but a lower correlation between numbers of tweets and retweets. For a direct comparison, we correlated just the repin rate with the number of pins, and found that the correlation between number of pins and repins was far lower ($\rho = 0.34$) than between the number of pins and followers, reported above.

Activity without Followers
We also examined the extreme end of our dataset: users who have no followers at all but nevertheless have activity on their pins. These users make up about 1.4% of our dataset, and strengthen the case for following being less than essential for content discovery; obviously, other users managed to find the content of these pinners, despite their having no followers.

Active Followers — RQ2
We now address RQ2, whether followers interact with their followees' content. Understanding the way users interact with the others they follow and those they don't is crucial to evaluating whether the follow mechanism is utilized for content discovery and sharing. It is well-known that only a small percentage of users' followers actively interact with them. This has been shown for Facebook, Twitter, and Cyworld (see the Related Work section). However, since following on Pinterest is supposed to be based on shared interests — a way to find new content to curate, and 'interactions' are mostly just a way of collecting content, we would expect users to receive interactions from a large percentage of their followers. We find, however, that this is not true: on average, only 12.3% of a user's followers have ever engaged in even a single action (like, comment, or repin) on any of the user's pins. The distributions of percentages of followers interacting for each user with any activity on the boards we crawled is shown in Figure 7a. Likes and comments once again have a more uneven distribution than repins; not only do those users who do interact comment and like much less than they repin, there are also *fewer* users who engage in either of the two activities than who repin. Even repins are only done by a small percentage of followers: only 19% of users have ever received even a single repin from more than 20% of their followers.

Content from Other Sources — RQ3
Unlike some SNSs, Pinterest is an open network — that is, users are not limited to interacting with their friends. Any user can like, comment on, or repin any other user's pins, without having to follow that user. This means that users can potentially find content for their 'collections' in places other than their home feeds or the boards of their followees. We wish to determine whether this capability is utilized by Pinterest users. Do they take content from users whom they don't follow? If so, how often? What percentage of activity does non-follower activity represent?

Figure 7b shows the distribution of the proportion of all repins, likes, and comments that are done by followers, for each user with at least one instance of the corresponding activity on their crawled boards. For the majority of users, only a small percentage of the interaction on their boards is from their followers (board or full-user) — the median is 24%. The remaining 76% are likes, comments, and repins done by non-followers. This is very different from Flickr, where between 47% and 71% (depending on photo visibility) of all interactions (comments) come from followers [42]. To confirm that these percentages were not skewed by a few prolific non-followers who engaged in a large amount of activity, we extracted the number of *unique* non-followers who interacted with each user's content. The median is 34, showing that most users have a significant number of non-followers interacting with their content; each non-follower contributes an average of just 1.4 actions. Non-following interacters make up a median of 88% of all unique users interacting with a user's pins.

These non-followers clearly discovered the content in this case somewhere other than their home feeds. (Since they are not following that user or board and therefore do not see the user's content in their home feed.) Yet they make up the vast majority of unique users interacting with the content. Together, these statistics strongly imply that users discover a great deal of the content they collect for themselves in places other than their home feeds, or their followers' boards.

Following as a result of content discovery — RQ4

Until now, we have discussed following as a driver of content discovery. In this section, we discuss following as a *result* of content discovery. As part of its mission to help users find content they like, Pinterest provides various mechanisms for content discovery besides the home feed. As we described previously, there are several pages containing a sampling of content from around the site. Among these are the Category pages, each with a sampling of pins from that category; and a Popular page, which hosts a continuous feed of pins with a large amount of activity. The pins remain on the boards of the users who posted them, but they can be interacted with from these 'sample' pages without visiting the original board; all activity done to a pin while on one of these pages is counted as part of the pin's activity as displayed on the original board. Users can browse these pages to find content and repin it to their own boards. They can also use the search feature to search for pins that match their keywords. Commensurate with its emphasis on following, Pinterest recommends that once a user finds and possibly interacts with new content, they should follow the board where the content was posted, or the entire pinner.[7] After all, if they liked this pin, they may like the other content on that board or pinned by that user as well. The idea that repins and likes on the Category or Popular pages are the gateway to new followers is frequently posited by Pinterest users and marketing experts [9].

We set out to investigate whether users do indeed go on to follow the users or boards whose pins they discover through search or the Category and Popular pages. The high proportion of non-follower activity that we reported in the previous section is the first indication we find that this behavior may not be very common. We assume that following that results from content discovery occurs within a fairly short time after the discovery and interaction. Thus, if the non-followers had indeed followed the user or board after interacting with the content, they would have been listed in our dataset as followers, since we crawled each board only once.[8] To further investigate, we compared the average number of actions contributed by each interacting follower of a user to the average number of interactions from each interacting non-follower. Each interacting non-follower contributes 1.4 interactions, on average, while the average interacting follower produces 3.4, almost almost two-and-a-half times as many (difference in means is significant, $p < .0001$). This difference suggests that non-followers, who have to find the board in some other way than viewing its pins in their home feed, seem to mostly find individual pins and do not interact with the rest of the pins on the board.

Being 'Popular' does not increase follower count

Lacking site usage data and search logs, we cannot trace the content that users found when searching. We do, however, have access to the Popular and Category pages. We therefore attempt to determine whether users who find interesting content on one of these pages go on to explore the rest of that pinner's content. We repeatedly visited the Popular page and collected a randomly-chosen single pin from the top row of pins (these are the most visible) as soon as it appeared on the page. We then crawled the board where each pin originated from,[9] once immediately and then again several hours later, after people had a chance to see the pin and possibly click through to, and follow, the board or user. We collected 1,013 Popular pins and found that, on average, each Popular pin had 673 likes, comments, and repins. (None had less than 111; 45% had over 500, and nearly a quarter had over 1,000.) In the vast majority of cases, the large amount of activity was clearly a result of being featured, since each popular pin had a median of 224 times more activity than an average pin on the same board. Despite all this activity, 60% of boards and 56% of users did not gain a single follower, and 80% gained 3 or fewer — a per-user average gain of one new follower for every 450 repins, likes, and comments — for those who gained at all. In addition, the majority of the 15% of users who did gain seven or more followers already had six or seven-digit numbers of followers; many are large retailers or successful bloggers with high visibility outside of Pinterest and multiple avenues for acquiring new followers. These factors mean that the additional followers they gained cannot be easily attributed to one of their pins being featured on the Popular page. Nor do users seem particularly interested in the other pins on the original board of a Popular pin; 70% of boards containing a Popular pin see no extra activity on their other pins after the pin is featured; those which do get some additional activity get a single repin, like, or comment on the non-featured pins for every 146 actions on the featured one.

It seems, then, that users do not go on to follow the 'out-of-network' users (or boards) whose content they discover and collect. After a pin is listed on the Popular page and receives high visibility and a large amount of additional activity, the number of followers of the user and board it belongs to barely change. We also found a wide disparity between the number of interactions on the Popular pins and the rest of the pins on the same board. This difference obviously cannot be accounted for by activity from the board's followers, since that

[7] http://help.pinterest.com/en/guide/following-and-your-home-feed

[8] Of course, there is always the small possibility that we managed to visit the page in the short time between when a user interacted with a pin and when they followed its pinner, but the chance of that happening is low enough that we assume that the number of pins affected, if any, is not significant.

[9] The first 100–300 pins; these are most likely to be looked at by a casual visitor.

should affect all pins. Nor can it be explained by differences in topic, since all pins on a board generally belong to the same category. It is likely that the additional activity came from non-followers who found the pin on the Popular page, but may not have even clicked through to the board. These users also do not appear to go on to follow the pinners who posted the content they enjoyed.

DISCUSSION

The results we present above show that the follow mechanism on Pinterest, and the resultant i.e. social browsing, does not seem to be very heavily utilized by users for content discovery and curation. The role of the home feed seems to be smaller than might be expected, given the site's emphasis on following as a driver of content discovery. We speculate that users see Pinterest as something of a 'mini-Web', an interlinked network of content that can be easily searched, browsed, and bookmarked.

Unlike Twitter and Facebook, for instance, Pinterest is not very time-bound; much of the content on it can still be relevant and engaging long after it was posted. The live-stream aspect of the home feed is therefore less essential. It is also (deliberately) centered around content, rather than people, which reduces the necessity for following: a status update about a friend's new baby is only interesting because you know the friend; a pin about a great recipe is useful even when it comes from a stranger. Pinterest may have hoped that this content-centrism would mean that people would follow others they didn't know solely for their content. What actually seems to be happening, however, is that people turn to other methods of content discovery, as discussed above. The many mechanisms that the site provides for finding desirable content means that users don't have to wait for their followees to post something interesting — they can look for it themselves.

Our findings have important implications for those who wish to promote goods or ideas on Pinterest. Rather than just seeking to maximize their number of followers, they will have to develop ways to engage with large numbers of users who have no interest in following them. These users may be very open, however, to repinning and otherwise engaging with their content.

Our results also add to the sparse existing knowledge on interest-based homophily on social network sites. While the example of Pinterest shows that it is possible to build a follow network based solely on common interests, it also shows that this network may not be robust enough to be heavily used by its members for content discovery when more attractive options are available.

CONCLUSION AND FUTURE WORK

In this paper, we study the effectiveness and utilization of the Pinterest follow mechanism for content discovery. Pinterest encourages users to form follow links on the site based on interest homophily, or similarity of interests. We sample user activity, in particular the curation and interaction actions of liking, commenting, and repinning, and examine it to determine whether or not users commonly use following and their home feeds as a way to find new content. We find that followers are not very active on the boards of their followees - on average, only 12.3% of a user's followers interact with any of the user's pins. Conversely, $> 75\%$ of the activity on an average user's boards is done by a large number of non-followers - 88% of all unique interacters on an average user's boards are non-followers. Users also do not seem to follow the boards or users that provide the content that they interact with, despite Pinterest's recommendation that they do so. Our results show that following seems to be a second-class mechanism for content discovery on Pinterest.

In this paper, we have mostly discussed characteristics of user interaction that can be extracted from the activity network. In the future, we would like to collect a larger, more fully connected dataset and examine various graph properties of the activity graph itself. We hope that this analysis will shed more light on Pinterest and the unique dynamics of user activity on a content-centered social network site. In particular, we are interested in the study of influence on Pinterest. Influence in social networks is an important topic of study because of the opportunities it presents for those who want to maximize the spread of ideas or advertising. Given the findings we present in this paper, we have reason to believe that the mechanisms of influence on Pinterest are both similar and different to those in other SNSs, and we would like to explore them in greater detail.

ACKNOWLEDGEMENTS

Bluma Gelley thanks Dan Pelleg, Mark Shovman, and Ido Guy for helpful discussions and valuable input. The authors thank Dorée Duncan Seligmann for her input in the preliminary stages of this work.

REFERENCES

1. Ahn, Y.-Y., et al. Analysis of topological characteristics of huge online social networking services. In *Proceedings of WWW 2007* (2007), 835–844.
2. Bakshy, E., et al. Everyone's an influencer: Quantifying influence on Twitter. In *Proceedings of WSDM 2011* (2011), 65–74.
3. Barabási, A.-L., and Albert, R. Emergence of scaling in random networks. *Science 286*, 5439 (1999), 509–512.
4. Bisgin, H., Agarwal, N., and Xu, X. Investigating homophily in online social networks. In *Proceedings of WI-IAT 2010*, vol. 1 (2010), 533–536.
5. Bonneau, J., et al. Eight friends are enough: Social graph approximation via public listings. In *Proceedings of WOSN 2009* (2009), 13–18.
6. boyd, d. Friendster and publicly articulated social networks. In *Proceedings of CHI 2004* (2004).
7. boyd, d., and Ellison, N. B. Social network sites: Definition, history, and scholarship. *Journal of Computer-Mediated Communication 13*, 1 (2007), 210–230.
8. Capocci, A., et al. Preferential attachment in the growth of social networks: The Internet encyclopedia Wikipedia. *Physical Review E 74*, 3 (2006), 036116.

9. Cario, J. E. *Pinterest Marketing: An Hour a Day*. John Wiley & Sons, 2012.

10. Cha, M., et al. Measuring user influence in Twitter: The million follower fallacy. In *Proceedings of ICWSM 2010*, vol. 10 (2010), 10–17.

11. Chang, S., et al. Specialization, homophily, and gender in a social curation site: Findings from Pinterest. In *Proceedings of CSCW 2014* (2014), 674–686.

12. Chen, W., Wang, C., and Wang, Y. Scalable influence maximization for prevalent viral marketing in large-scale social networks. In *Proceedings KDD 2010* (2010), 1029–1038.

13. Chen, W., Wang, Y., and Yang, S. Efficient influence maximization in social networks. In *Proceedings of KDD 2009* (2009), 199–208.

14. Choi, J. H.-j. Living in Cyworld: Contextualising Cy-ties in South Korea. In *Uses of Blogs*, A. Bruns and J. Jacobs, Eds. Peter Lang, New York, 2006, 173–186.

15. Chun, H., et al. Comparison of online social relations in volume vs interaction: A case study of Cyworld. In *Proceedings of IMC 2008* (2008), 57–70.

16. Cole, J., and Lebo, H. Surveying the digital future report. *USC Annenberg School Center for the Digital Future* (2013).

17. Constine, J. Pinterest hits 10 million u.s. monthly uniques faster than any standalone site ever -comscore. `http://techcrunch.com/2012/02/07/pinterest-monthly-uniques/`, 2012.

18. Donath, J., and boyd, d. Public displays of connection. *BT technology Journal 22*, 4 (2004), 71–82.

19. D'Onfro, J. This chart shows why pinterest is way more important than twitter or reddit for traffic referrals. `http://www.businessinsider.com/pinterest-advertising-big-deal-2014-4`, 2014.

20. Duggan, M., and Smith, A. Social media update 2013. *Pew Internet and American Life Project* (December 30, 2013).

21. Ellison, N. B., and boyd, d. Sociality through social network sites. *The Oxford handbook of Internet studies* (2013), 151.

22. Ellison, N. B., Steinfield, C., and Lampe, C. The benefits of Facebook"friends": Social capital and college students use of online social network sites. *Journal of Computer-Mediated Communication 12*, 4 (2007), 1143–1168.

23. Ellison, N. B., Steinfield, C., and Lampe, C. Connection strategies: Social capital implications of Facebook-enabled communication practices. *New Media & Society* (2011), 1461444810385389.

24. Feng, Z., et al. An empirical study of user behaviors on Pinterest social network. In *Proceedings of WI-IAT 2013*, vol. 1 (2013), 402–409.

25. Gilbert, E., et al. I need to try this?: A statistical overview of Pinterest. In *Proceedings of CHI 2013* (2013), 2427–2436.

26. Gilbert, E., and Karahalios, K. Predicting tie strength with social media. In *Proceedings of CHI 2009* (2009), 211–220.

27. Golder, S. A., and Yardi, S. Structural predictors of tie formation in Twitter: Transitivity and mutuality. In *Proceedings of SocialCom 2010* (2010), 88–95.

28. Goodman, L. A. Snowball sampling. *The Annals of Mathematical Statistics 32*, 1 (1961), 148–170.

29. Hall, C., and Zarro, M. Social curation on the website Pinterest.com. *Proceedings of the American Society for Information Science and Technology 49*, 1 (2012), 1–9.

30. Han, J., et al. Collecting, organizing, and sharing pins in Pinterest: Interest-driven or social-driven? In *Proceedings of SIGMETRICS 2014* (2014).

31. Henry, A. How Facebook is using you to annoy your friends (and how to stop it). `http://lifehacker.com/5987248/how-facebook-is-using-you-to-annoy-your-friends-and-how-to-stop-it`, 2013.

32. Hodkinson, P. Subcultural blogging? online journals and group involvement among UK Goths. In *Uses of blogs*, A. Bruns and J. Jacobs, Eds. Peter Lang, New York, 2006, 187–198.

33. Huberman, B., Romero, D. M., and Wu, F. Social networks that matter: Twitter under the microscope. *First Monday 14*, 1 (2008).

34. Hutto, C., Yardi, S., and Gilbert, E. A longitudinal study of follow predictors on Twitter. In *Proceedings of CHI 2013* (2013).

35. Incite, N. Friends & frenemies: Why we add and remove Facebook friends. *State of Social Media Survey* (2011).

36. Kamath, K. Y., Popescu, A.-M., and Caverlee, J. Board recommendation in Pinterest. In *Proceedings of UMAP Workshops 2013* (2013).

37. Kivran-Swaine, F., Govindan, P., and Naaman, M. The impact of network structure on breaking ties in online social networks: Unfollowing on Twitter. In *Proceedings of CHI 2011* (2011), 1101–1104.

38. Kwak, H., Moon, S. B., and Lee, W. More of a receiver than a giver: Why do people unfollow in Twitter? In *Proceedings of ICWSM 2012* (2012).

39. Lampe, C., Ellison, N., and Steinfield, C. A Face (book) in the crowd: Social searching vs. social browsing. In *Proceedings of CSCW 2006* (2006), 167–170.

40. Lauw, H. W., et al. Homophily in the digital world: A LiveJournal case study. *Internet Computing, IEEE 14*, 2 (2010), 15–23.

41. Lenhart, A., and Madden, M. Teens, privacy, and online social networks. *Pew Internet and American Life Project* (April 18,2007).

42. Lerman, K., and Jones, L. A. Social browsing on Flickr. In *Proceedings of ICWSM 2007* (2007).

43. Macskassy, S. A., and Michelson, M. Why do people retweet? Anti-homophily wins the day! In *Proceedings of ICWSM 2011* (2011).

44. Marlow, C., et al. Maintained relationships on Facebook. `http://overstated.net/2009/03/09/maintained-relationships-on-facebook`, 2009.

45. McPherson, M., Smith-Lovin, L., and Cook, J. M. Birds of a feather: Homophily in social networks. *Annual review of sociology* (2001), 415–444.

46. Miller, A. D., and Edwards, W. K. Give and take: A study of consumer photo-sharing culture and practice. In *Proceedings of CHI 2007* (2007), 347–356.

47. Minott, K. Is Pinterest an aspirational social network? `http://www.komverse.com/2013/04/04/is-pinterest-an-aspirational-social-network/`, 2013.

48. Mittal, S., et al. The pin-bang theory: Discovering the Pinterest world. *arXiv preprint arXiv:1307.4952* (2013).

49. Moore, R. J. Pinners be pinnin: How to justify Pinterests 3.8b valuation. `http://blog.rjmetrics.com/2014/05/07/pinners-be-pinnin-how-to-justify-pinterests-3-8b-valuation/`, 2014.

50. Motamedi, R., et al. Characterizing group-level user behavior in major online social networks. `http://onrg.cs.uoregon.edu/pub/CIS-TR-2013-09.pdf`, 2014.

51. Naaman, M., Boase, J., and Lai, C.-H. Is it really about me?: Message content in social awareness streams. In *Proceedings of CSCW 2010* (2010), 189–192.

52. Nogueira Cortimiglia, M., Ghezzi, A., and Renga, F. Social applications: Revenue models, delivery channels, and critical success factors-an exploratory study and evidence from the Spanish-speaking market. *Journal of theoretical and applied electronic commerce research 6*, 2 (2011), 108–122.

53. Ottoni, R., et al. Ladies first: Analyzing gender roles and behaviors in Pinterest. *Proceedings of ICWSM 2013* (2013), 457–465.

54. Quercia, D., Bodaghi, M., and Crowcroft, J. Losing friends on Facebook. In *Proceedings of WebSci 2012* (2012), 251–254.

55. report, S. Geolocation and activity analysis of Pinterest accounts by Semiocast. `http://semiocast.com/en/publications/2013_07_10_Pinterest_has_70_million_users`, 2013.

56. Romero, D. M., et al. Influence and passivity in social media. In *Machine learning and knowledge discovery in databases*. Springer, 2011, 18–33.

57. Romero, D. M., and Kleinberg, J. M. The directed closure process in hybrid social-information networks, with an analysis of link formation on Twitter. In *Proceedings of ICWSM 2010* (2010).

58. Seitzinger, J. Curate me! exploring online identity through social curation in networked learning. In *Proceedings of NLC 2014* (2014), 412–419.

59. Suh, B., et al. Want to be retweeted? Large scale analytics on factors impacting retweet in Twitter network. In *Proceedings of SocialCom 2010* (2010), 177–184.

60. Ugander, J., et al. The anatomy of the Facebook social graph. *arXiv preprint arXiv:1111.4503* (2011).

61. Van House, N. A. Flickr and public image-sharing: Distant closeness and photo exhibition. In *CHI'07 extended abstracts* (2007), 2717–2722.

62. Ver Steeg, G., and Galstyan, A. Information transfer in social media. In *Proceedings of WWW 2012* (2012), 509–518.

63. Viswanath, B., et al. On the evolution of user interaction in Facebook. In *Proceedings of WOSN 2009* (2009), 37–42.

64. Wilson, C., et al. User interactions in social networks and their implications. In *Proceedings of EuroSys 2009* (2009), 205–218.

65. Wirtz, B. W., Schilke, O., and Ullrich, S. Strategic development of business models: implications of the Web 2.0 for creating value on the Internet. *Long Range Planning 43*, 2 (2010), 272–290.

66. Wu, F., Wilkinson, D. M., and Huberman, B. A. Feedback loops of attention in peer production. In *Proceedings of CSE 2009*, vol. 4 (2009), 409–415.

67. Xiang, R., Neville, J., and Rogati, M. Modeling relationship strength in online social networks. In *Proceedings of WWW 2010* (2010), 981–990.

68. Zarro, M., Hall, C., and Forte, A. Wedding dresses and wanted criminals: Pinterest.com as an infrastructure for repository building. In *Proceedings of ICWSM 2013* (2013), 650–658.