

Social Bootstrapping: How Pinterest and Last.fm Social Communities Benefit by Borrowing Links from Facebook

Changtao Zhong¹, Mostafa Salehi², Sunil Shah³, Marius Cobzarenco⁴,
Nishanth Sastry¹, Meeyoung Cha⁵

¹King's College London ²University of Tehran ³UC Berkeley ⁴Last.fm ⁵KAIST
¹{changtao.zhong, nishanth.sastry}@kcl.ac.uk, ²mostafa_salehi@ut.ac.ir, ³sunil.shah@berkeley.edu,
⁴marius@last.fm, ⁵meeyoungcha@kaist.edu

ABSTRACT

How does one develop a new online community that is highly engaging to each user and promotes social interaction? A number of websites offer friend-finding features that help users bootstrap social networks on the website by copying links from an established network like Facebook or Twitter. This paper quantifies the extent to which such *social bootstrapping* is effective in enhancing a social experience of the website. First, we develop a stylised analytical model that suggests that copying tends to produce a giant connected component (i.e., a connected community) quickly and preserves properties such as reciprocity and clustering, up to a linear multiplicative factor. Second, we use data from two websites, Pinterest and Last.fm, to empirically compare the subgraph of links copied from Facebook to links created natively. We find that the copied subgraph has a giant component, higher reciprocity and clustering, and confirm that the copied connections see higher social interactions. However, the need for copying diminishes as users become more active and influential. Such users tend to create links natively on the website, to users who are more similar to them than their Facebook friends. Our findings give new insights into understanding how bootstrapping from established social networks can help engage new users by enhancing social interactivity.

Categories and Subject Descriptors

H.3.5 [Online Information Services]: Commercial Services, Data Sharing, Web-based services; J.4 [Social and Behavioral Sciences]: Sociology

Keywords

Social Bootstrapping; Friend Finder Tools; Community Design; Social Property; Social Interaction; Copied Networks

1. INTRODUCTION

How to design online communities and maintain users participation is a fundamental problem for website designers. Many websites now try to incorporate a social networking aspect to enhance

user engagement and create active communities. Making a website “social” typically involves linking users together and providing some kind of awareness of the linked users’ activities to each other. Studies have found that such social networking aspects facilitate community formation in learning [1, 6, 14], working [7, 17], medicine [9] and online games [5, 8] applications.

However, in creating such a social experience on a website, designers face an important choice: should they create an entirely new social network embedded within the site? Or should they instead connect users who are already linked together on an established social network such as Facebook and Twitter? The latter option has recently become a possibility, with both Facebook and Twitter opening up their social graphs to third-party websites, who can write friend-finder tools that help users select and import friendship links from these established networks into their own service (e.g., through the open graph protocol [10]).

We term this act of copying existing friends from an established social network onto a third-party website as *social bootstrapping*. Social bootstrapping has direct implications on how a new online social network community can grow quickly. However, this problem is complex to examine with real data because it involves user interaction across multiple heterogeneous networks. To this end, we gather massive amounts of data from Facebook, Pinterest, and Last.fm involving tens of millions of nodes and billions of links and explore the potential benefits and limitations of social bootstrapping¹. We seek to evaluate how such bootstrapping could affect the user community and to what extent copying links contributes to social structure and user engagement as the new website matures.

Although copying clearly enriches the number of social links on the new website, it is not *a priori* clear whether links borrowed from a general-purpose social network such as Facebook would be appropriate for content-driven sites, which typically attempt to link users interested in similar content. Additionally, social bootstrapping involves a two-step process. First, to copy a Facebook friendship, for example, both users joined by the friend link have to independently decide to connect their accounts on the third-party website with their Facebook accounts. Next, they have to select which of their friends to import into the third-party website and choose this particular link. Thus, social bootstrapping can be limited in the number of links that get copied over.

Nonetheless, social bootstrapping becomes effective for user engagement and community formation if it creates the sort of structure conducive to social interaction and increased user activity on the third-party website. Therefore, using a combination of analytical models and empirical studies, we focus on three important

Copyright is held by the International World Wide Web Conference Committee (IW3C2). IW3C2 reserves the right to provide a hyperlink to the author’s site if the Material is used in electronic media.
WWW’14, April 7–11, 2014, Seoul, Korea.
ACM 978-1-4503-2744-2/14/04.
<http://dx.doi.org/10.1145/2566486.2568031>.

¹The Pinterest datasets used in this paper are shared for wider community use at http://www.inf.kcl.ac.uk/staff/nrs/projects/cd-gain/social_bootstrapping.html.

structural properties, namely connectivity, reciprocity, and clustering. We study the effect of copying on these properties and how social interactions are affected in turn. We also study how copying evolves among more active and influential users and build up a picture of the importance of creating native links on the new website.

We first develop a stylized model of copying as a process of sampling links from the established network. To mimic the two-step filtering process described above, we propose the Link Bootstrapping Sampling (LBS) model as a variation of induced subgraph sampling [15]. Under this simple analytical model, we study the emergence of a giant connected component in the copied network. We demonstrate that when copying from a typical network with a heavy-tailed degree distribution, a giant component emerges even with a small amount of sampling, which suggests that social bootstrapping may be an effective means of increasing user engagement and creating a connected community.

We then empirically study copying, using data from two large websites, Pinterest and Last.fm, which include friend-finding tools to copy friends from Facebook. We make efforts to tease out various social effects. We study the structural properties of the copied subgraph, comparing it to the subgraph of links created natively on the website, and find that copying enriches reciprocity and clustering of the local structure. Both reciprocity and clustering are shown to be important for social interactions, indicating that social bootstrapping successfully promotes user engagement.

However, copying links yields diminishing returns. As users become more active and influential on the new website, they create proportionally more native links than copied ones. Native links offer a benefit over copied links: users connecting natively on Pinterest and Last.fm tend to be more similar to each other in their tastes than with the ones copied from Facebook. This is an important observation for long-term user engagement, as prolific users tend to engage more with native links and fine-tune the local relationships to meet their interests. As a result, we conclude that while “copying” links is essential to bootstrap one’s network, the opposite “weaning” process is equally important for long lasting user engagement.

To the best of our knowledge, this paper is the first to demonstrate how content-driven websites like Pinterest and Last.fm can benefit from social bootstrapping by copying links from established networks like Facebook. Through extensive analysis of cross-network data, we are able to describe both how new social communities are seeded by social bootstrapping, and how users grow beyond the bootstrapped links to create strong communities natively. We believe our findings have strong implications for the design of new content-driven web communities.

2. A LINK BOOTSTRAPPING MODEL

In this section, we propose a simple analytical model of social bootstrapping to gain insight about its implications on network structure. Our model allows us to analytically examine how copying links affects key structural features that facilitate social interactions in the target network, such as reciprocity, clustering and the formation of a Giant Connected Component (GCC).

2.1 Terminology

Social bootstrapping refers to the act of copying existing friend links from a source social network onto a third-party website to create a target social network. We define several sub-networks below to describe this phenomenon:

Source network: The social graph of an established social network like Facebook (Fb for short), which contains a signif-

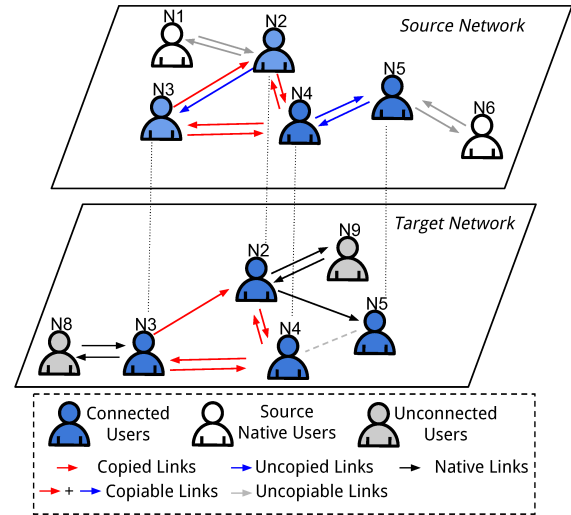


Figure 1: The structure of social bootstrapping.

icant number of nodes and links (e.g., 1.19 billion monthly active users as of 2014²). The source network is displayed as the upper layer in the toy example in Fig. 1. Note that some users, such as $N1$ and $N6$ are *source native* and are present only in the source network.

Target network: The relatively new third-party network that allows users to copy links from established networks, displayed as the lower layer in Fig. 1. *Connected nodes* are the subset of all nodes in target network that have used the “Friend Finder” tool to connect their accounts to the source network. In the toy example, blue nodes, i.e., $N2$, $N3$, $N4$ and $N5$ are the connected nodes. Grey nodes, i.e., $N8$, $N9$ are *unconnected nodes* who either exist only on the target network or have chosen not to connect their accounts on the source network to their identity on the target network. Within the target network, social links copied from the source network are called *copied links* and those created natively are called *native links*. Copied links in the target network may be directed even if they are copied from the undirected source network. Copied links are a subset of *copiable links*, the set of all links between connected nodes in the source network. We take Pinterest (*Pnt* for short) and Last.fm (*Lfm* for short) as two target networks of interest.

Copied network: The social subgraph of the target network solely containing copied links and all *connected nodes*. In Fig. 1, the copied network contains the red edges and all blue nodes. We call the network copied from Facebook as *Fb-copied*.

Native network: The subgraph of the target network that only contains native links and the corresponding nodes at either end of each native link. In the toy example, the native network is the subgraph made up by black edges and nodes linked by them, i.e., $N2$, $N3$, $N5$, $N8$ and $N9$. Nodes can be in copied and native networks at the same time, but links are either copied or native. We call the native networks for Pinterest and Last.fm *pnt-native* and *lfm-native*, respectively.

2.2 A random bootstrapping process

We now introduce a simple analytical model that represents copying a social link as a simplified random sampling process. We

²<http://newsroom.fb.com/Key-Facts>

propose a two-step model, called the *Link Bootstrapping Sampling (LBS)*, which is a variation of the induced subgraph sampling process [15]. In the first step, users of the target network have to self-select to connect their accounts on the target network with the source network. In the second step, users have to select which of their friends from the source network to import onto the target network. Under this stylized model, we obtain expressions for the resulting degree distributions of the copied network and a condition for the emergence of a giant connected component in that network. Although the model considers directed target networks, it can be trivially adapted for undirected networks.

Formally, let $G = (V, E)$ be the graph representing the globally-established source network, where V is the set of nodes, and E is the set of links between pairs of nodes. We assume that the LBS copying process randomly samples each node N_i with a probability $p_1(N_i)$. We call this the *node sampling rate*. Each selected node independently selects each of its neighbours N_j in the original network with a probability $p_2(N_i, N_j)$. This is the *link sampling rate*. Let $S \subset V$ be a sample of nodes, and $L \subset E$ denote the collection of sampled links to be found in G among subset S . Then, the subgraph $G(S) = (S, L)$ represents a copied network. A (directed) copied link from N_i to N_j appears in $G(S) = (S, L)$ if both events (node and link sampling) happen, i.e., with probability $p_1(\cdot)p_2(\cdot, \cdot)$. We call this the *link copy probability*.

2.3 Giant component in the copied subgraph

We assume for ease of exposition that the node and link sampling rates p_1 and p_2 are uniform across all nodes and edges respectively to set the following link copy probability

$$p_e = p_1 p_2 \quad (1)$$

However, similar results on the appearance of GCC, Eq. (4), can also be obtained using mean field approximations for p_e . This allows us to consider more realistic assumptions such as the node sampling rate being proportional to the degree to reflect the possibility that more social nodes connect their source and target accounts with higher probability. Alternatively, we may also consider conditions such as the link sampling rate being proportional to the number of common friends between the nodes on either side of the link to reflect the possibility that socially closer links are copied over with higher probability.

The probability that a node will have exactly $k^{i(o)}$ links in the copied subgraph $G(S)$ is given by:

$$p_c^{i(o)}(k^{i(o)}) = \sum_{k_0^{i(o)}=k^{i(o)}}^{\infty} p_s(k_0^{i(o)}) \times \binom{k_0^{i(o)}}{k^{i(o)}} (p_e)^{k^{i(o)}} (1-p_e)^{k_0^{i(o)}-k^{i(o)}} \quad (2)$$

where $p_s^{i(o)}(k)$ is the in- (out-) degree distribution of the source network G . Similarly, the joint degree distribution of obtaining a node with in-degree j and out-degree k in the copied subgraph is:

$$p_c(j, k) = \sum_{j_0=j}^{\infty} \sum_{k_0=k}^{\infty} p_s(j_0, k_0) \binom{j_0}{j} (p_e)^j (1-p_e)^{j_0-j} \times \binom{k_0}{k} (p_e)^k (1-p_e)^{k_0-k} \quad (3)$$

where $p_s(j, k)$ is the joint degree distribution of source network G .

We are now ready to state an important property, relating the link copy probability to the emergence of a giant component:

PROPERTY 1. *When copying from an undirected source network (or equivalently, every link occurs in both directions), a giant component appears in the copied network if*

$$p_e \geq \frac{\langle k \rangle'}{\langle k^2 \rangle'}, \quad (4)$$

where averages are computed with respect to the degree distribution in the source network.

PROOF. In a directed network of arbitrary degree distribution $p(j, k)$, a GCC exists if [22]:

$$\langle jk \rangle \geq \langle k \rangle \quad (5)$$

For the copied network induced by the Link Bootstrapping Sampling, the average of the degree and joint degree distribution of the subgraph $G(S)$ sampled by the Link Bootstrapping Sampling method from a network G with joint distribution $p_s(j, k)$, can be calculated, respectively, as $\langle jk \rangle = p_e^2 \langle jk \rangle'$ and $\langle k \rangle = p_e \langle k \rangle'$, where $\langle jk \rangle'$ and $\langle k \rangle'$ are computed with respect to the original joint degree distribution (i.e., $p_s(j, k)$). Thus, we can rewrite Eq. (5) as:

$$p_e \geq \frac{\langle k \rangle'}{\langle jk \rangle'} \quad (6)$$

Since the source network is undirected, in- and out-degrees are completely correlated. i.e., $\langle jk \rangle' = \langle k^2 \rangle'$. With this, Eq. (6) reduces to Eq. (4). \square

Thus the link copy probability p_e at which a GCC emerges in the sampled subgraph under Link Bootstrapping Sampling depends in a simple and intuitive manner on the properties of the source network. For instance, if the source network can be modeled as an Erdos-Rényi random network, where the degree distribution is given by a Poisson distribution with parameter λ , i.e. $p_s(k) = \frac{\lambda^k \exp(-\lambda)}{k!}$, the first and second moments are given by $\langle k \rangle' = \lambda$ and $\langle k^2 \rangle' = \lambda^2 + \lambda$. Thus, the link copy probability must exceed $p_e > \frac{1}{\lambda+1}$ for a GCC to exist. If the source network is scale-free, a GCC emerges easily if the degree distribution of the source network has infinite second moments $\langle k^2 \rangle'$. In particular, we know that in a power law distribution $p_s(k) = ck^{-\gamma_s}$, all moments of order $m > \gamma_s - 1$ are infinite. Thus, if the Link Bootstrapping Sampling is applied to copy links from an undirected *scale-free network* with power law exponent $\gamma_s < 3$, a GCC will come into existence even with very low link copy probability ($p_e \rightarrow 0$). In general, the larger the second moment, or equivalently, the larger the variation in the degree distribution, the easier it is (i.e., the lower the link copy probability needed) for a giant component to emerge.

2.4 Other properties

Next, we study the effect of Link Bootstrapping Sampling (using uniform link and node sampling rates) on two other properties in the copied subgraph, which are thought to be correlated with social interaction: reciprocity and clustering coefficient. Both increase proportionally with the link sampling rate.

2.4.1 Reciprocity

First, we study the effect of bootstrapping on R^c , the reciprocity of the copied network, defined as the proportion of links which exist in both directions, among all copied links. We have

$$R^c = 1 - \frac{2m_s[p_2(1-p_2)]}{2m_s p_2} = p_2 \quad (7)$$

where m_s is the number of links in the source network. Thus, reciprocity is defined by the link sampling rate; higher link sampling rates result in higher expected reciprocity.

2.4.2 Clustering

Next, we obtain an expression for the clustering coefficient of the copied network. Taking the copied network to be an uncorrelated undirected network with arbitrary degree distribution, the clustering coefficient takes the value [23]:

$$C^c = \frac{1}{n} \frac{[\langle k^2 \rangle - \langle k \rangle]^2}{\langle k \rangle^3} \quad (8)$$

where $\langle k \rangle$ and $\langle k^2 \rangle$ are the first and second moments of the degree distribution, respectively in the copied network, and n is the number of sampled nodes. Writing these moments in terms of the corresponding moments $\langle k \rangle'$, $\langle k^2 \rangle'$ of the source network [24]:

$$\langle k \rangle = p_e \langle k \rangle', \quad \langle k^2 \rangle = p_e^2 \langle k^2 \rangle' + p_e(1 - p_e) \langle k \rangle' \quad (9)$$

Substituting these formulae into Eq. (8), we have

$$C^c = \frac{1}{n} \frac{p_e^4 [\langle k^2 \rangle' - \langle k \rangle']^2}{p_e^3 \langle k \rangle'^3} = \frac{1}{n} p_e N C^s = \frac{p_1 N}{n} p_2 C^s = p_2 C^s \quad (10)$$

where C^s is the clustering coefficient of the source network.

The Link Bootstrapping Sampling with uniform node and edge sampling preserves the clustering coefficient of the source network, up to a multiplicative factor corresponding to the link sampling probability. This means that copying links from a source network that has a high level of clustering results in a copied network also with a proportionally high level of clustering.

3. DATASETS

Having gained initial insight into copying from an analytical perspective, in the rest of the paper, we take an empirical approach and examine social bootstrapping using datasets from two very different websites, Pinterest and Last.fm. These are considered as target networks in our analysis. In both cases, we study copying from Facebook as the source network. Our datasets include extensive social graphs from both target websites, as well as corresponding graphs from Facebook. In addition, it includes nearly all activities from selected periods on both websites. The data that we collected from Pinterest is shared to the research community, while much of the Last.fm data is already available through a public API.

3.1 Pinterest

Pinterest is a photo sharing website that allows users to store and categorise images. Images added on Pinterest are termed *pins* and can be created in two ways. The first way is *pinning*, which imports from a URL external to pinterest.com. A second is *repinning* from an existing pin. All pins are organised into *pinboards* or boards, which belong to one of 32 globally specified categories. In addition to pinning or repinning, users can like a pin or comment on a pin.

The social graph of Pinterest is created through users *following* other users or boards they find interesting. We call social links created in this way native links. In addition to this method, users are able to connect with their Facebook and Twitter accounts and import their social links into Pinterest. The *Find Friends* function provides a list of Facebook and Twitter friends who are also registered on Pinterest. Users can select some of them to follow on the Pinterest website, which we call *copied links*.

Table 1 summarizes the Pinterest dataset, consisting of the social graph on Pinterest, the corresponding nodes and edges on Facebook, and activities on the Pinterest site in January 2013. To obtain the Pinterest social graph, we used a snowball sampling technique, starting to crawl from a seed set of 1.6 million users which we

collected in advance. In total 68.7 million Pinterest users and 3.8 billion directed edges between them were obtained. For each user, we checked whether there was a connected Facebook account, and gathered basic profile information such as gender and profile, as well as basic statistics such as the number of pins, likes, followers, and followees. Of the 68.7 million, 40.4 million were Facebook-connected users, who have 2.4 billion links between them on Pinterest.

We next separate the 2.4 billion edges into those which are present on Facebook (i.e., are Fb-copied), and those which are native to Pinterest (Pnt-native). To identify the Fb-copied portion of the network, we used the Facebook API to individually check whether a Pinterest link between two connected users was also present between the corresponding Facebook accounts³. We find that 0.98 billion links between connected users are also on Facebook. These form our Fb-copied network. Pnt-native links were identified by excluding the Fb-copied network from our Pnt network.

In a previous study [27], we had collected nearly all activities within Pinterest, during the period from January 3rd to 21st, 2013. The crawl proceeded in two steps: first, to discover new pins, we visited each of the 32 category pages once every 5 minutes, and collected the latest pins of that category. Then, for every pin obtained this way, we visited the webpage of the pin every 10 minutes. A pin's webpage lists the 10 latest repins and the 24 latest likes⁴; we added these to our dataset, along with the approximate time of repins, likes and comments (if any). Through these regular visits, we captured almost all the activity during our observation period. We estimate that the fraction of visits which resulted in missed activities stands at 5.7×10^{-6} for repins and 9.4×10^{-7} for likes. In total, 8.5 million users (termed as *active users*), 38.0 million repins and 19.9 million likes were included.

Amongst these active users, there are 5.2 million connected to Facebook. We crawled the Facebook pages of these 5.2 million connected active users, and attempted to obtain their Facebook friend lists. Due to privacy settings, only 2.3 million users' social links could be obtained. Together, this collection of Facebook edges constitutes a subgraph of 444.2 million edges (Table 1c). Of these, 141.9 million are *copiable* links, i.e., edges between connected users who are on both Facebook and Pinterest.

	Nodes	Links
Pnt network	68,665,590	3,871,570,784
Fb-copied	40,472,339	983,520,986

(a) Target social graph

	Timespan	Repins	Likes
Activities	03-21 Jan 2013	38,041,368	19,907,874

(b) Activities in Pinterest

Nodes	links
2,322,473	444,216,279

(c) Facebook network

Table 1: The social graph among all Pinterest users.

³Note that checking whether a pair of users are friends is affected by users' privacy setting. That is, it is unknown for us whether two users are friends or not if both of them had set their friend lists as private. Also, we assume that a link which exists both on Facebook and the target networks is a copied link, first made on Facebook and then copied to the target network. Although we expect this to be the case normally, it is possible for user pairs to link to each other separately on Facebook and Pinterest, or link first on Pinterest, and subsequently on Facebook. We are unable to distinguish these cases from links copied using friend finder tools.

⁴This setting has been changed in April 2013.

3.2 Last.fm

Last.fm is a music discovery and recommendation website. Users can log what they listen to using a multitude of applications which support a variety of different operating systems and audio playback devices. This activity is known as *scrobbling*. Scrobbled data is used to provide recommendations to users via collaborative filtering methods and is displayed publicly on users' profile pages. Users can *love* tracks, a mechanism akin to liking a pin on Pinterest. These tracks are also displayed on their profile page and can optionally be shared to Facebook.

Last.fm offers a social network in which users can friend each other. A friendship between users can be considered as a bidirectional link, similar to that which Facebook offers. Friends of each user are displayed on their profile page, and when logged in, users are shown what their friends have scrobbled and what tracks their friends have loved.

Users can connect their Facebook accounts to Last.fm in three ways. The first is using their Facebook account to bootstrap basic profile information when they first sign up. Second, Last.fm offers a friend finder tool which connects to third party services such as Facebook, Google Mail and Yahoo! to look for contacts on those services, who also use Last.fm. Note that the Facebook friend finder can only find other friends who have already connected their Facebook account to Last.fm. The third and most recent method is that users who share event attendance and loved tracks via their Facebook profile connect the two accounts as a result.

We considered a subset of the overall Last.fm user base by looking only at a sample of 1.8 million users who had, at some point in their history on Last.fm, connected a Facebook account to their Last.fm account using one of the methods. For each consenting user, we had access to their Last.fm social graph, basic profile information and their Facebook username. Of these users, 904,132 users use the Last.fm social features (i.e., have friendship edges in the Last.fm social network). Between these users, we extracted a subgraph of 12.3 million directed edges (or 6.15 million friendships) which forms the Lfm network.

For each of these 12.3 million Last.fm edges, we checked whether the friendship is also present in Facebook, using the Facebook API. Through this procedure, we identified 2.8 million copied edges between 600,000 users. Privacy settings meant that we were unable to validate friendships for approximately 200,000 users. We identify the Lfm-native network by eliminating copied edges.

We measure these users' activities on Last.fm in two ways: First, we measure listening activity by counting their scrobbles. Second, we measure site usage from their website access log. Both measures cover the period Jan 1–Jun 22, 2013.

Finally, we extract friend request data for requests sent during 2012 between Facebook connected users who were active on the site during that period (defined as those who have visited the site over 100 times during 2012). This data includes who initiated the friend request, as well as how it was made (i.e., through the friend finder tool, or natively on Last.fm) and whether the request was accepted, ignored, cancelled or is still pending. This contains about 141,000 users and 1.1 million friend requests.

	Nodes	Links
Lfm network	1,843,020	12,291,658
Fb-copied	592,992	2,787,000

Table 2: Social graph of Last.fm users

4. STRUCTURAL BENEFITS OF COPYING

Section 2 showed that copying can produce desirable properties such as a giant component, reciprocity and clustering. We now

empirically analyse how these structural properties of the copied network compare to the natively created network. We examine implications of the differences we find, for social interaction in the target social network community.

4.1 Copied network has higher reciprocity

Reciprocity is known to indicate positive bidirectional interaction between a pair of users, which is also known to increase user longevity in the system [11, 2, 28]. Here, we attempt to examine the effect of copying on creating structurally stronger bidirectional social ties, by defining *reciprocity ratio* as the fraction of social links that are *reciprocal*, or bidirectional. For a node in a network, let her follower (or following) set in the target network (e.g., Pnt or Lfm) be *ind* (or *out*) and her friend set copied from the source network (e.g., Fb) be *fr*. Then the reciprocity ratios of that user in the entire target networks, and its partition into Fb-copied, and native networks are as follows:

$$R_{copied} = \frac{|fr \cap ind \cap out|}{|fr \cap (ind \cup out)|},$$

$$R_{native} = \frac{|(ind - fr) \cap (out - fr)|}{|(ind - fr) \cup (out - fr)|}.$$

On some services like Pinterest, users *follow* others unilaterally, creating directional links. We study the extent to which *follow* acts are reciprocated and become bidirectional. On other services like Last.fm, users initiate a *friend* request, which needs to be approved by the other party before a friendship link is instantiated, creating bidirectional links by default. In this case, we induce a directional network by using data about historical friendship requests, treating the initial friend request as a *follow*, and examine the extent to which such requests are approved by the other party, creating reified bidirectional friendship links.

Fig. 2a shows that in both Pinterest and Last.fm, the reciprocity ratio is higher in links which are also found on Facebook, than on natively created links. Although in some cases, a link copied in one direction could be reciprocated by the other party merely in order to be “social” or “polite”, the link creation creates an opportunity for social interaction on the target website, and reciprocity could promote positive bidirectional social interactions (We verify this in §4.4). The figure also indicates that the reciprocity ratio for Last.fm is significantly higher than Pinterest. This is consistent with user studies in previous work [27] which found that Last.fm users easily accept requests. Fig. 3 shows that copying is extremely important for establishing reciprocal relationships. In Pinterest, a large proportion of users' reciprocal links are in fact those copied from Facebook. In Last.fm it is slightly different: the fraction is relatively smaller than Pinterest, which we think is again because users tend to accept requests easily.

4.2 Copied network shows higher clustering

Next we explore the impact of copying on another popular measure of a strong social structure, clustering or the degree to which users share common friends. Fig. 2b shows that in both websites, users have much higher clustering co-efficients on the copied network than on the network natively created on the website. Thus copying not only promotes reciprocal social interactions, but also creates a much denser social network structure in the target website.

4.3 Copying enhances connectivity

The increased clustering and reciprocity are properties relating to local structure around a node. Copied links are also crucial for connectivity, a global (network-wide) property. Fig. 2c confirms that both the Pinterest and Last.fm copied networks have a giant

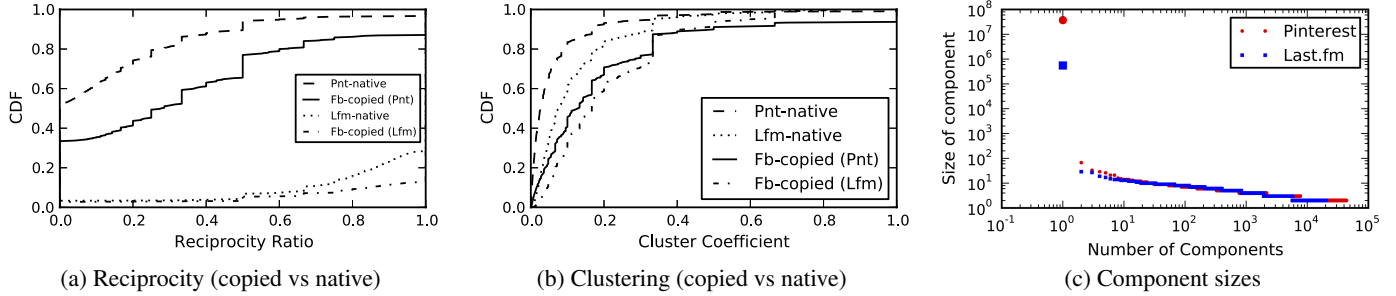


Figure 2: Properties of copied subgraph. (a) CDF of per-user fraction of links reciprocated in copied and natively created networks. More links are reciprocated in the copied network. (b) Per-user CDF of clustering coefficients in natively created and copied subgraphs of Pinterest and Last.fm (0 valued-points not shown). Clustering coefficients are higher in the copied network. (c) Distribution of the sizes of connected components on the FB-Copied network in the Pinterest and Last.fm datasets.

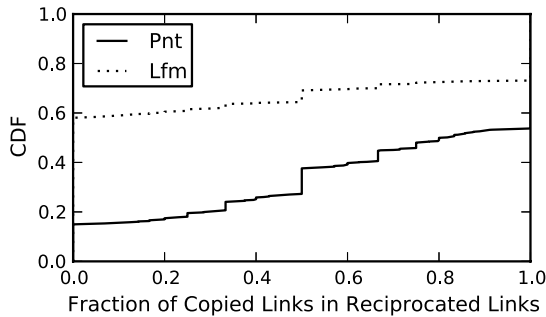


Figure 3: CDF of per-user fractions of Fb-copied links among reciprocated links in target networks. Many users have high proportions of Fb-copied links implying that copied links are important for establishing bidirectional or reciprocated relationships.

component. The largest component comprises 0.91 (Pinterest) and 0.93 (Last.fm) of all the connected nodes (i.e., nodes present on both source and target networks). Furthermore, this component encompasses 0.53 (Pinterest) and 0.66 (Last.fm) of all the nodes in the corresponding target network.

4.4 Implications for social interactions

So far, we have shown that copying links results in a higher level of reciprocity and clustering, representing a stronger and denser social structure than its low-clustering and low-reciprocity native counterpart. While these properties are *expected* to improve social interaction [18, 26], we ask whether the benefits of these structural properties are seen in the social interactions of the target network.

In order to determine the benefits of a close-knit structure, we examine one of the most popular activities on the Pinterest network, repinning. Our main question is whether copied links promote higher levels of repins. To measure this effect, we first define the concept of a *social repin*, which is a repin in which a user repins a pin of someone whom she follows. We then define the *social repin network*, as the subgraph of links in the Pinterest network over which at least one social repin happens in our data.

We examine how the social repin network selectively samples the underlying network of Pinterest. First, we ask what proportion of a user’s reciprocated and directed (unreciprocated) links have incurred repins. Fig. 4a shows that repins happen more easily over

reciprocated links. Next, in Fig. 4b we compare the clustering coefficient of users in the social repin network to the clustering coefficient of the underlying graph. Users have significantly higher clustering coefficient when we remove the links over which no repins happen. This suggests that social interactions tend to be directed towards the closer friends of a user, within highly clustered communities.

These results show that the social repin graph is richer in reciprocated links and is more highly clustered than the underlying network. Since reciprocal links and high clustering nodes will have more social repins, it is straightforward to infer that the copied network, which is higher in both reciprocity and clustering coefficient, should promote more social repins. This is proved by Fig. 4c, which shows that a larger fraction of social repiners tend to be from the copied network than from the natively created network.

5. WEANING FROM FACEBOOK

While copying links provides instant bootstrapping advantages by incurring a close-knit local structure (i.e., high reciprocity and clustering), there is a limit to which a user can copy links from Facebook. Beyond a certain point, a user may no longer find other Facebook friends to copy over. It is natural to ask whether this creates engagement bottlenecks for users as they become more prolific on the target network, or whether they find alternative solutions.

In this section, we describe a collective “weaning” process, through which users move away from their reliance on Facebook copied links to building new relationships natively on target websites. We find that users, as they become more active and influential on Pinterest and Last.fm, establish more native links within these services and copy less from Facebook. We discuss why users “go native” in this way and suggest a possible cause: through native links, users may find others similar to themselves on the target website.

5.1 Measures of activity and influence

To quantify the level of user activity on Pinterest, we employ three different measures: the numbers of boards created, pins made (including repins of other users’ pins), and likes of others’ pins. The level of influence of a user is merely the activity of *other* users directed towards that user, i.e., the number of repins and likes received by that user for her pins. In the case of Last.fm, while there is no direct measure of social influence, we have two measures of activity: the number of scrobbles and number of hits to the website.

The results of this section are robust in the sense that they all hold for each measure of activity and influence defined above. Due

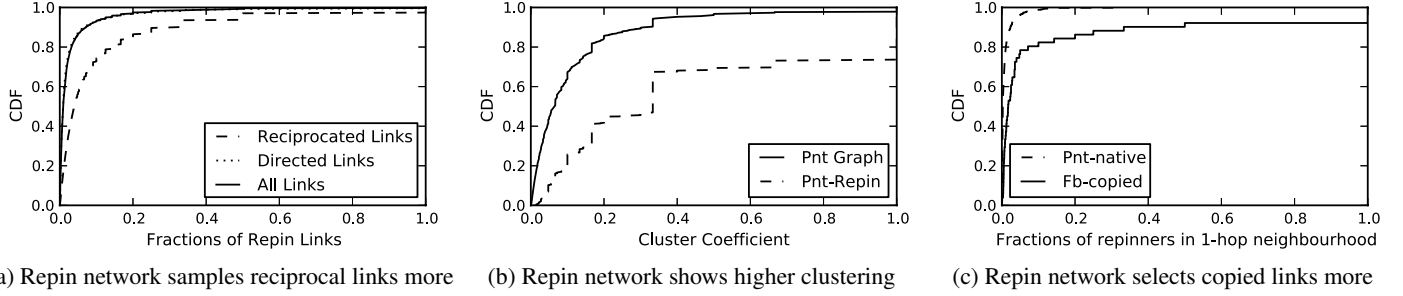


Figure 4: How the social repin network samples the Pinterest graph (0-valued points not shown): (a) CDF of fraction of users' reciprocated and unreciprocated (directed) links, which are included in the repin network. A greater fraction of reciprocated links than directed links have repin activity. (b) CDF of users' clustering coefficients in the Pinterest graph and the repin network. The repin network has higher clustering, indicating that users' social repins are directed more at closer friends. (c) CDF of the fractions of users' natively created and copied (Fb-copied) links which are sampled by the repin network. Copied links tend to have more repins.

to space limitations, however, results are selectively shown for only some of these measures.

5.2 Active and influential users copy fewer links

In order to study levels of copying, we introduce a measure called the *copy ratio*. Denoting the set of all friends in the target network as *all* and the friend set copied from the source network (i.e., Facebook) as *fr*, the copy ratio in a undirected network, such as Last.fm's, is defined as:

$$CR = \frac{|all \cap fr|}{|all|}$$

For a directed network, representing a node's follower (resp., following) set in the target network (i.e., Pinterest) by *ind* (resp., *out*), we define the *follower copy ratio* and *following copy ratio* as:

$$CR_{ind} = \frac{|ind \cap fr|}{|ind|}$$

$$CR_{out} = \frac{|out \cap fr|}{|out|}$$

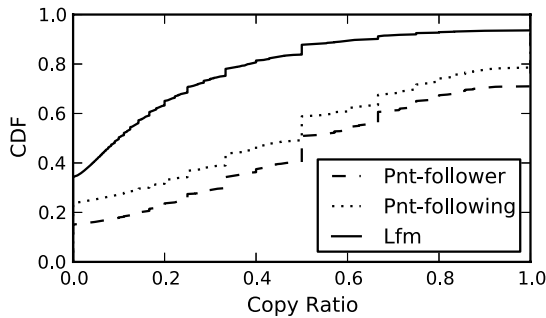


Figure 5: CDF of copy ratio for connected users.

Fig. 5 shows the cumulative distribution of copy ratios (CR). Based on the trend shown, we may divide users into three groups. Approximately 20% of users are *Pinterest natives* (resp., *Last.fm natives*), who only create links natively on the website (i.e., CR=0). A second category (20–30%) link only to their Facebook friends

(i.e., CR=1) and can be termed *Facebook expats*. The majority (50–60%), however, are *bi-networked*, relying on a mixture of both native and copied links ($0 < CR < 1$). The Last.fm network contains a larger fraction of natives who have zero copy ratio (CR=0) and a smaller fraction of Facebook expats (CR=1) compared to Pinterest.

Next, we compare the activity of these three categories of users in Fig. 6a–c. All combinations of activity measures and the two copy ratios show that Facebook expats (CR=1) whose social links are entirely copied from Facebook are the least active, whereas bi-networked users ($0 < CR < 1$) with a mixture of native and copied links are the most active. Pinterest and Last.fm natives (CR=0) who do not copy at all are in the mid range. This implies that users who start with the bootstrapping advantage tend to move away from reliance on the existing Facebook network and start building new links natively (hence becoming bi-networked) as they become active members on Pinterest or Last.fm.

Fig. 6d–e drill down further and examine how the copy ratios change as activity levels increase, for the case of pins and scrobbles in Pinterest and Last.fm respectively. This demonstrates a clear inverse relationship between the activity levels and copy ratio, with users who pin or scrobble a lot tending to have lower levels of copying—that is, higher activity levels are associated with lower copy ratio. Fig. 6f shows that this result extends to measures of influence on Pinterest. We find that users who are influential, measured by repins, tend to have lower copy ratios. (In the case of Last.fm, we omit this analysis due to the lack of an influence measure.) Overall, the results above indicate that as users settle down on the new service and become more active and influential, their investment in natively formed links increases proportionally.

5.3 Influential and active users remain social, but with native rather than copied friends

Next we take a deeper look at the relationship between the increase in activity or influence level of a user and his level of social interaction. In order to quantify the level of social interaction, we again define the concept of a *social repin* as a repin where the user who is repinning follows the original pinner. We define a user's *social repin ratio* for activity (or influence) as the fraction of social repins made (or received) among all repins made (or received).

Fig. 7a–b shows that users who are more active (or influential) tend to make (or receive) proportionally more social repins in relation to their activity (or influence) level, *confirming that social*

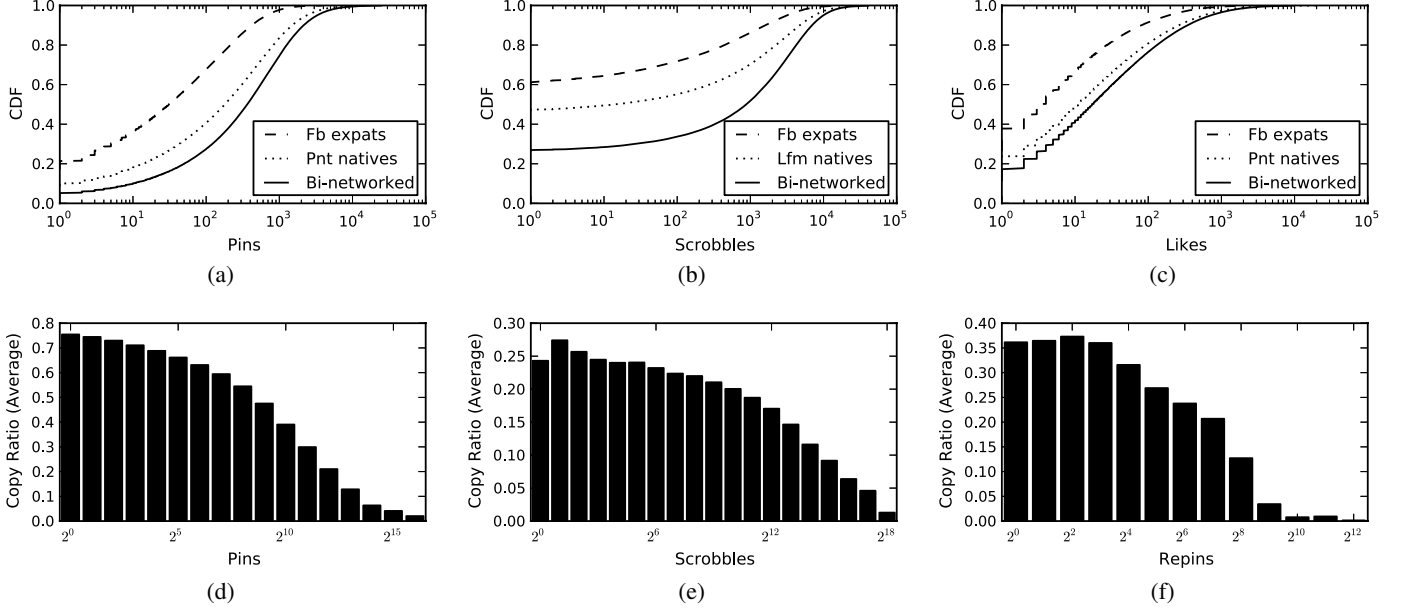


Figure 6: (a-c) Facebook-expats (CR=1) exhibit lesser activity in terms of pins, likes and scrobles, compared to Pinterest and Last.fm natives (CR=0) and bi-networked users ($0 < CR < 1$), who are the most active. (d-e) Higher activity levels measured by pins (scrobles) are associated with lower following copy ratio in Pinterest (Last.fm). (f) Users who are influential on Pinterest, as measured by repins, tend to have lower copy ratios.

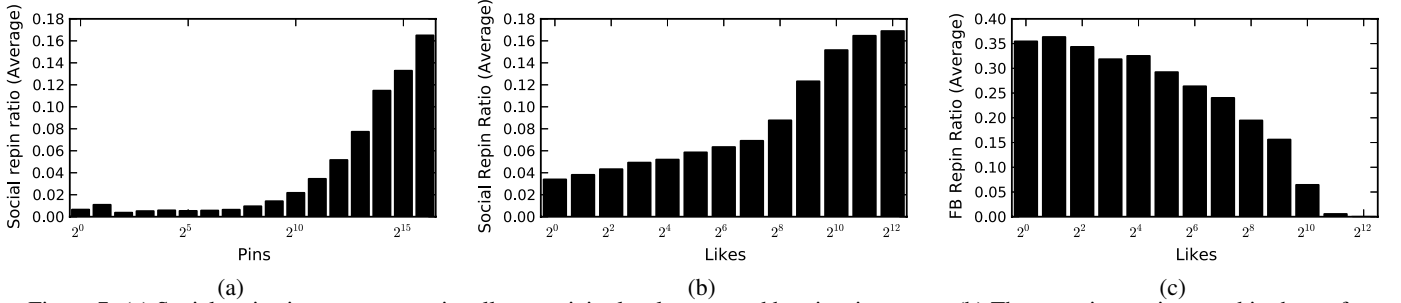


Figure 7: (a) Social repins increase proportionally as activity level, measured by pins, increases. (b) The same increasing trend is shown for influence, measured by likes received. (c) However, the fraction of social repins over links copied from Facebook decreases proportionally as influence, measured by likes received, increases.

interaction continues to be increasingly essential for active (or influential) users.

We also focus on social repins and ask whether copied links promote social repins. We define the *Facebook repin ratio* for activity (or influence) as the fraction of social repins made (or received) over Fb-copied links among all social repins made (or received). Fig. 7c reveals that as activity (or influence) levels increase, social repins happening over copied links decrease.

5.4 Weaning, biases and community evolution

We conclude by asking how the nature of the target network community would evolve as users ‘wean’ from copying to make more native links. To understand this, we study user preferences or biases in the kinds of links they copy and the links they make natively. We also seek to understand the role that copying plays in creating more native links.

User studies in previous work on Pinterest [27] identified that Pinterest users most value the social aspect of the service that helps

them find people with similar tastes in pictures. Therefore, we examine whether natively created links on Pinterest enables discovery of individuals with a more similar taste than those with copied links. Specifically, if I_1 is the set of user u_1 ’s board categories, and I_2 is the set of u_2 ’s, we define their similarity as $s = \frac{|I_1 \cap I_2|}{|I_1 \cup I_2|}$. Fig. 8a confirms that according to this measure, users connected by native Pinterest links are more similar to each other than those connected by Facebook-copied links.

Fig. 8a also shows that there is no difference in similarity between users who are copied and users who are not copied over from Facebook. This implies that users are not selecting Facebook friends to copy based on similarity. On Last.fm, native links also show higher similarity than Facebook copied links, which in turn show higher similarity than Facebook links which were deliberately not copied. We conjecture that on some websites like Pinterest, there might be social norms at play which are being used to decide which links could be copied over.

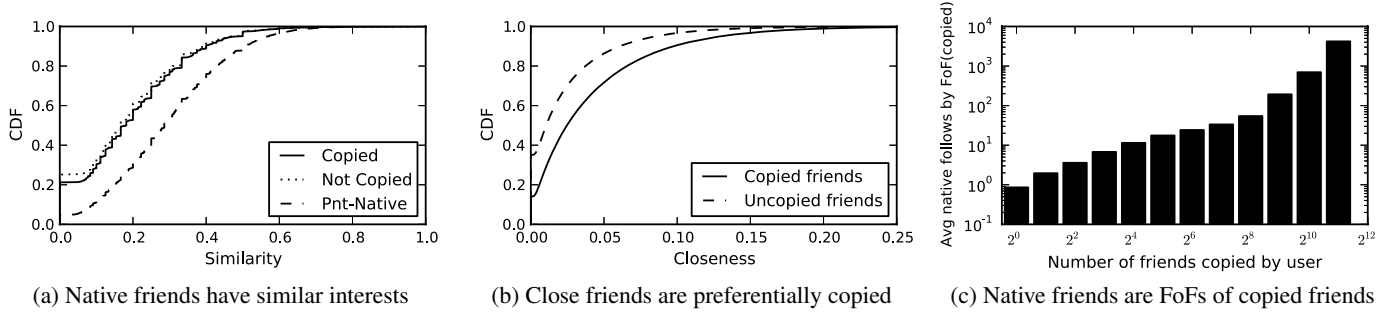


Figure 8: User choice in copying links (Pinterest data) (a) CDF of similarity between users linked by copied, uncopied and pnt-native links, showing that native friends are more similar to a user than copied friends, but copied and uncopied friends do not differ significantly in tastes. (b) Per-user CDF of closeness between copied and uncopied friends. It shows that copied friends are closer than uncopied friends. (c) The more friends a user copies and follows in the target network, the more follows she gets, from exposure to friends of the copied friends.

In Fig. 8b, we study whether closeness of friends has a role in deciding which friends to copy. In our analysis, we use the similarity of users’ friend lists to show their closeness: if A ’s friend list is L_A and B ’s is L_B , we say their closeness is $\frac{L_A \cap L_B}{L_A \cup L_B}$. Fig. 8b shows that closeness between copied friends is higher than between uncopied friends. In Last.fm, however, there is no significant difference between closeness of copied and uncopied friends.

Together these results suggest that Pinterest users tend to use the “friend finder” tool to copy close friends they know from established source networks like Facebook, but when they discover new friends on the target network, they tend to prefer users with similar tastes. Thus, as native links become more important and numerous than copied links, we expect the target network to become more interest-based. By contrast, Last.fm users tend to prefer copying and natively linking to users who share similar music tastes. Thus, while there is no universal pattern for how users on different target networks copy links, it appears that in both cases, the links, and hence the target communities, tend to become more interest-based over time.

However, copying continues to be important for the creation of native links over which interaction happens, even in networks like Pinterest, where copying appears to be governed by norms of social closeness – Fig. 8c examines links over which social repin interactions happen over a sample representative day, and shows that users who have copied more of their friends from the source to target network, tend to have more native followers who are *friends of her friends* on the target network. i.e., copying creates the opportunity for users in the immediate social community of nodes in the copied sub-graph to discover and follow them, creating new native links, over which social interaction happens.

6. RELATED WORK

Many websites try to incorporate a social networking aspect to enhance user engagement and community interaction. Social networks are known to facilitate the formation of learning communities, foster student engagement and reflection, and enhance overall user experience for students in synchronous and asynchronous learning environments [1]. Social networks are also the core of the design of new user-driven communities around health issues [9] and have been utilised to facilitate community formation in environments ranging from professional settings [17] to online games [5, 8]. Including the above studies, most existing research on the formation and evolution of online social network communities has focused on single networks. In contrast, this paper evaluated interac-

tions between two *different* networks: a generic social network and a target network on the content-driven website.

Multilayer networks (or also called multiplex, heterogeneous, interdependent, or multi-relational networks in literature) describe the fact that users may belong to different social networks (or layers) at the same time in real world. Each network layer could have particular features different from the others. A number of studies have looked into multilayer networks, including modelling of the formation and evolution of multilayer networks based on preferential attachment models [19, 20, 21]. In particular, resilience of cooperative behaviours is known to be enhanced by a multilayer structure [13] and in some cases, cascading failures may occur in interacting networks [3]. Multilayer structures are also known to speed up diffusion in networks [12].

We proposed a copying process for networks as a model called Link Bootstrapping Sampling (LBS). This model can be seen as a variation of Induced Subgraph Sampling (ISS) [15], which randomly selects a subset of nodes and observes all links between selected nodes. Compared to ISS, the LBS model introduces an additional link sampling step: each selected node further selects a *subset* of its links for observation. Then, based on this model, we derived conditions for the emergence of a giant connected component in copied networks. Lee *et al.* [16] also examine the emergence of a giant connected component in multi-layer networks. However their study is restricted to the Erdős-Rényi model and seeks to answer a different question of how the correlation of node degrees in different layers affects the emergence of a giant component.

Another aspect of this paper was the large-scale empirical study across two different networks. Empirical analysis of multiple networks is relatively uncommon. Szell [25] collected data from an online game and extracted networks of six different types of one-to-one interactions between the players. Then, both reciprocity and clustering were studied for each layer of the network. In contrast, our dataset shows the process of copying links between two independent websites, the source and target, where the original purpose of the link in the source network may be quite different from the intended purpose for the copied link in the target network.

Finally, this paper is related to the series of studies that investigate the motivation of users in creating social network links. One study found that professionals use internal social networking to build stronger bonds with their weak ties and to reach out to employees they do not know [7]. Another study identified that social links have high predictive power in determining which newcomers will continue to engage with the service in the future [4]. Other

studies have demonstrated that properties such as reciprocity and clustering promote interaction in natively created social graphs [18, 26]. We have shown that such positive effects of social links also apply to links copied from unrelated *external* social networks.

7. CONCLUSIONS

This paper studied the impact of social bootstrapping—the act of copying one’s social ties or links from a source network to a target network. This is a popular practice enabled by many new social network services and has implications on how a new online social network community can grow quickly. We gathered massive amounts of data from Facebook, Pinterest, and Last.fm involving tens of millions of nodes and billions of links to understand this new phenomenon. We proposed a simple analytical model and used it to gain insight into the social bootstrapping process. Among a number of findings, we highlight that a “copying” process is useful to initiate social interaction in the target network, as one may expect. However, a “weaning” process, where a user moves away from copied social links and builds social relationships natively in the new network is essential for longer lasting user engagement. To the best of our knowledge, this paper is the first to utilize large-scale cross network data in understanding the interplay between heterogeneous services in terms of bootstrapping a network and engaging users to form a cohesive, interacting community.

Acknowledgements

This research was partially supported by the UK Engineering and Physical Sciences Research Council via Grant No. EP/K024914/1 and the Basic Science Research Program through the National Research Foundation of Korea funded by the Ministry of Science, ICT & Future Planning (2011-0012988).

8. REFERENCES

- [1] D. E. Baird and M. Fisher. Neomillennial user experience design strategies: Utilizing social networking media to support “always on” learning styles. *Journal of educational technology systems*, 34:5–32, 2005.
- [2] M. Boguna and M. A. Serrano. Generalized percolation in random directed networks. *Physical Review E*, 72(1), July 2005.
- [3] S. V. Buldyrev, R. Parshani, G. Paul, H. E. Stanley, and S. Havlin. Catastrophic cascade of failures in interdependent networks. *Nature*, 2010.
- [4] M. Burke, C. Marlow, and T. Lento. Feed me: motivating newcomer contribution in social network sites. In *CHI*. ACM, 2009.
- [5] D. Choi and J. Kim. Why people continue to play online games: In search of critical design factors to increase customer loyalty to online contents. *CyberPsychology & behavior*, pages 11–24, 2004.
- [6] G. Conole and J. Culver. The design of cloudworks: Applying social networking practice to foster the exchange of learning and teaching ideas and designs. *Computers & Education*, 54:679–692, 2010.
- [7] J. DiMicco, D. R. Millen, W. Geyer, C. Dugan, B. Brownholtz, and M. Muller. Motivations for social networking at work. In *CSCW*. ACM, 2008.
- [8] N. Ducheneaut and R. J. Moore. The social side of gaming: a study of interaction patterns in a massively multiplayer online game. In *CSCW*, pages 360–369. ACM, 2004.
- [9] G. Eysenbach. Medicine 2.0: Social networking, collaboration, participation, apomediation, and openness. *Journal of Medical Internet Research*, 10, Aug. 2008.
- [10] Facebook. Facebook developers’ webpage on open graph. Last accessed, 28 Jun 2013.
- [11] E. Fehr and S. Gächter. Fairness and Retaliation: The Economics of Reciprocity. *The Journal of Economic Perspectives*, 2000.
- [12] S. Gómez, A. D. Guíler, G. G. nes, C. J. P. Vicente, Y. Moreno, and A. Arenas. Diffusion Dynamics on Multiplex Networks. *Physical Review Letters*, 2013.
- [13] J. Gómez-Gardeñes, I. Reinares, A. Arenas, and L. M. M. Flórida. Evolution of cooperation in multiplex networks. *Scientific reports*, 2012.
- [14] G. Heiberger and R. Harper. Have you facebooked astin lately? using technology to increase student involvement. *New Directions for Student Services*, 2008(124):19–35, 2008.
- [15] E. D. Kolaczyk. *Statistical Analysis of Network Data: Methods and Models*. Springer Publishing Company, Incorporated, 1st edition, 2009.
- [16] K.-M. Lee, J. Y. Kim, W.-k. Cho, K.-I. Goh, and I.-M. Kim. Correlated multiplexity and connectivity of multiplex random networks. *New Journal of Physics*, 2012.
- [17] T. Y. Lee, C. Dugan, W. Geyer, T. Ratchford, J. Rasmussen, N. S. Shami, and S. Lupushor. Experiments on motivational feedback for crowdsourced workers. In *ICWSM*, 2013.
- [18] S. A. Macskassy. On the study of social interactions in twitter. In *ICWSM*, 2012.
- [19] M. Magnani and L. Rossi. Formation of multiple networks. In *International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction*, 2013.
- [20] V. Nicosia, G. Bianconi, V. Latora, and M. Barthelemy. Growing multiplex networks. *CoRR*, 2013.
- [21] B. Podobnik, D. Horvatić, M. Dickison, and H. E. Stanley. Preferential attachment in the interaction between dynamically generated interdependent networks. *Europhysics Letters*, 2012.
- [22] N. Schwartz, R. Cohen, B. D. Avraham, A. L. Barabási, and S. Havlin. Percolation in directed scale-free networks. *Physical Review E*, 2002.
- [23] M. Serrano and M. Boguñá. Clustering in complex networks. I. General formalism. *Physical Review E*, 2006.
- [24] S. W. Son, C. Christensen, G. Bizhani, D. V. Foster, P. Grassberger, and M. Paczuski. Sampling properties of directed networks. *Physical Review E*, 2012.
- [25] M. Szell, R. Lambiotte, and S. Thurner. Multirelational organization of large-scale social networks in an online world. *Proceedings of the National Academy of Sciences*, 107:13636–13641, Aug. 2010.
- [26] C.-Y. Teng and Lada A. Adamic. Longevity in second life. In *ICWSM*, 2010.
- [27] C. Zhong, S. Shah, K. Sundaravadivelan, and N. Sastry. Sharing the loves: Understanding the how and why of online content curation. In *ICWSM*, 2013.
- [28] Y.-X. Zhu, X.-G. Zhang, G.-Q. Sun, M. Tang, T. Zhou, and Z.-K. Zhang. Influence of reciprocal links in social networks. arXiv e-print 1306.5538, June 2013.