Evolution of Ego-networks
in Social Media with Link Recommendations

Luca Maria Aiello
Nokia Bell Labs
Cambridge, United Kingdom
luca.aiello@nokia-bell-labs.com

Nicola Barbieri
Tumblr
New York, NY, USA
barbieri@yahoo-inc.com

ABSTRACT

Ego-networks are fundamental structures in social graphs, yet the process of their evolution is still widely unexplored. In an online context, a key question is how link recommender systems may skew the growth of these networks, possibly restraining diversity. To shed light on this matter, we analyze the complete temporal evolution of 170M ego-networks extracted from Flickr and Tumblr, comparing links that are created spontaneously with those that have been algorithmically recommended. We find that the evolution of ego-networks is bursty, community-driven, and characterized by subsequent phases of explosive diameter increase, slight shrinking, and stabilization. Recommendations favor popular and well-connected nodes, limiting the diameter expansion. With a matching experiment aimed at detecting causal relationships from observational data, we find that the bias introduced by the recommendations fosters global diversity in the process of neighbor selection. Last, with two link prediction experiments, we show how insights from our analysis can be used to improve the effectiveness of social recommender systems.

Keywords

Network evolution, ego-networks, link recommendation, groups, communities, social media, Tumblr, Flickr

1. INTRODUCTION

Ego-centric social networks (ego-networks) map the interactions that occur between the social contacts of individual people. Because they provide the view of the social world from a personal perspective, these structures are fundamental information blocks to understand how individual behaviour is linked to group life and societal dynamics. Despite the growing availability of interaction data from online social media, little research has been conducted to unveil the structure and evolutionary dynamics of ego-networks [4, 30]. In an online context, people expand their social circles also as a result of automatic recommendations that are offered

to them, which makes it harder to disentangle spontaneous user behavior from algorithmically-induced actions.

We aim to provide an all-round description about how ego-networks are formed and how automated contact recommendations might bias their growth. We do so by analyzing the full longitudinal traces of 170M ego-networks from Flickr and Tumblr (§3), answering several open research questions about the shape of their boundaries, their community structure, and the process of neighbor selection in time (§4). The richness of the data we study allows for the identification of those Tumblr links that have been created as a result of recommendations served by the platform, which positions us in an unique standpoint to investigate the impact of link recommender systems on the process of network growth.

Some of our key findings are:

- The backbone of a typical ego-network is shaped within the initial month of a node's activity and within the first ∼50-100 links created. In that period, new contacts are added in larger batches and the main communities emerge. Unlike global social networks, whose diameter shrinks in time, the average distance between nodes in ego-networks expands rapidly and then stabilizes.

- The selection criteria of new neighbors change as new contacts are added, with popular contacts being more frequently followed in earlier stages of the ego's life, and friends-of-friends being selected in later stages. The neighbor selection is also heavily driven by the ego-network's community structure, as people tend to grow different subgroups sequentially, with an in-depth exploration strategy.

- The link recommender system skews the process of ego-network construction towards more popular contacts but at the same time restraining the growth of its diameter, compared to spontaneous behavior. With a matching experiment aimed at detecting causal relationships from observational data, we find that the bias introduced by the recommendations fosters diversity: people exposed to recommendations end up creating pools of contacts that are more different from each other compared to those who were not exposed.

The outcomes of our analysis have theoretical implications in network science and find direct application in link recommendation and prediction tasks. We run a prediction experiment (§5) to show that simple temporal signals could be crucial features to improve link prediction performance, as the criteria of ego-network expansion vary as the ego grows older. In a second experiment, we test the algorithmic capability to tell apart spontaneous links from recommendation-

induced ones. This ability opens up the way to train link predictors that mitigate existing algorithmic biases by suggesting links whose properties better adhere to the natural criteria that people follow when connecting to others.

2. RELATED WORK

Structure and dynamics of social networks. For decades, network science research has explored extensively the structural and evolutionary properties of online social graphs and of the communities they encompass [24, 11, 31, 7, 41, 39, 54, 55, 49], unveiling universal patterns of their dynamics. Individual connectivity and activity are broadly distributed [39, 51]; the creation of new links is driven by reciprocation, preferential attachment [38], triangle closure [33], and homophily [2, 58]. Globally, the number of edges in a social network grows superlinearly with its number of nodes, and the average path length shrinks with the addition of new nodes [34], after an initial expansion phase [1]. The regular patterns that drive the link creation process have enabled the development of accurate methods for link prediction and recommendation [28] based on either local [35] or global structural information [9, 8, 45]. Fine-grained temporal traces of user activity in online social platforms opened up new avenues to investigate in detail the impact of time on network growth [60]. For example, the relationship between the node age and its connectivity has been measured in several online social graphs including Flickr [33, 57].

Ego-networks. To date, not much research has been conducted on how nodes build their *local* social neighborhoods in time. Research done by Aranboldi et al. has looked into ego-networks of online-mediated relationships including the Facebook friendship network [4] and the Twitter follow graph [5, 3], as well as professional relationships such as Google Scholar's co-authorship network [6]. Using community detection, hierarchical clusters are discovered, in agreement with Robin Dunbar's theory on the hierarchical arrangement of social ego-circles [59]. Similar findings have been confirmed by independent studies on the Facebook network [19]. In the attempt of comparing the properties of the global network with those of ego-networks, recent studies found that local structural attributes are characterized by local biases [27] that are direct implications of the friendship paradox [21]. Multiple techniques have been proposed to discover social or topical sub-groups within ego-networks [53, 37, 40, 13], but with little attention to the dynamics of their growth. Kikas et al. conducted one of the few studies touching upon the the temporal evolution of ego-networks, using a dataset of Skype contacts [30]. They find that most edges are added in short bursts separated by long inactivity intervals.

Effect of social recommender systems. In the past years, computer scientists developed increasingly effective contact recommender systems for online social media [26]. Only recently, the community has adopted a more critical standpoint with respect to the *effects* that those recommendations may have on the collective user dynamics. Algorithms based on network proximity are better suited to find contacts that are already known by the user, whereas algorithms based on similarity of user-generated content are stronger at discovering new friends [16]. Surveys administered to members of corporate social networks revealed that contact recommendations with high number of common
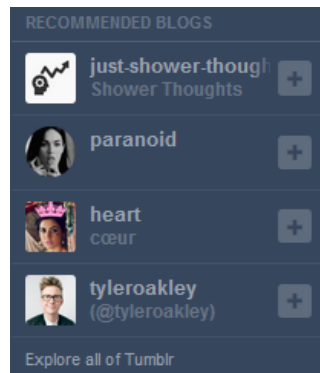

Figure 1: Tumblr's contact recommender system.

neighbors are usually well-received [17]. Recommendation-induced link creations have a substantial effect on the growth of the social graph; for example, the introduction of the "people you may know" service in Facebook increased considerably the number of links created and the ratio of triangle closures [60]. A recent study on Twitter compared the link creation activity before and after the introduction of the "Who To Follow" service [48], showing that popular nodes are those who most benefit from recommendations. On a wider perspective, the debate around recommender systems fostering or limiting access to novel information is still open. On one hand, recommenders may originate a filter bubble effect by providing information that increasingly reinforces existing viewpoints [42]. On the other hand, recent research has pointed out that individual choices, more than the effect of algorithms, limit exposure to cross-cutting content [10]. It has been argued that recommender systems have a limited effect in influencing people's free will. Observational studies on Amazon found that 75% of click-throughs on recommended products would likely have occurred also in the absence of recommendations [44]. In the context of link recommendation systems, a key open question is how they affect ego-network diversity.

3. DATASET AND PRELIMINARIES

We study two social media platforms that differ in both scope and usage. The data includes only interactions between users who voluntarily opted-in for research studies. All the analysis has been performed in aggregate and on anonymized data.

3.1 Tumblr

Tumblr is a popular social blogging platform. The types of user-generated content range from simple textual messages to multimedia advertising campaigns [15, 25]. Users might own multiple blogs, but for the purpose of this study we consider blogs as users, and we will use the two terms interchangeably. Users receive updates from the blogs they follow; the following relationship is directional and might not be reciprocated. In Tumblr, 326 million blogs and 143 billion posts have been created[1] since the release of the platform to the public in July 2007. We extracted a large random sample of the social network in October 2015, which includes

---

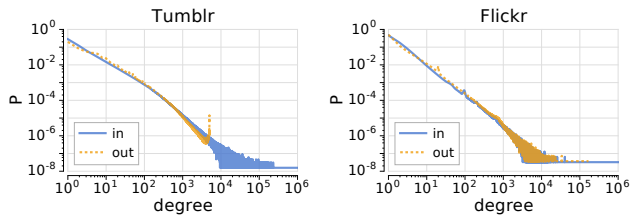[1]https://www.tumblr.com/about (Dec 2016)

Figure 2: Degree distributions in Tumblr ($\mu_{in} = 108$, $\mu_{out} = 58$) and Flickr ($\mu_{in} = 19$, $\mu_{out} = 21$).

almost $7B$ follow links created between $130M$ public blogs over approximately 8 years. All the social links are marked with the exact timestamps of their creation, which allows for a fine-grained longitudinal analysis of the network evolution.

In October 2012, Tumblr launched a new version of its *recommended blogs* feature. On the web interface, a short-list of four recommended blogs is displayed in a panel next to the user's feed (Figure 1). Users can get more recommendations by clicking on "explore". At every page refresh, the shortlist may change according to a randomized reshuffling strategy that surfaces new recommended contacts from the larger pool. Tumblr's link recommendation algorithm is not publicly disclosed, but it considers a mixture of two signals: topical preferences and network structure. The user's tastes are estimated since the onboarding phase, in which registrants are asked to indicate their preference on a set of pre-determined topics organized in a taxonomy (e.g., sports, football). The topical profile helps to overcome the cold-start problem. As the number of contacts grows, new blogs are recommended following the triangle closure (friend-of-a-friend) principle.

For all the links created after January 2015, we can reliably estimate if they have been created as an effect of a recommendation. This information is inferred by combining the log of recommendation impressions (i.e., when recommendations are visualized by the user) with the log of link creations. When the link is created shortly after the recommendation is displayed, we count the link creation as triggered by a recommendation.

3.2 Flickr

Flickr is a popular photo-sharing platform in which users can upload a large amount (up to 1 TB) of pictures and share them with friends. Users can establish directed social links by following other users and get updates on their activity. Since its release in February 2004, the platform has gathered almost 90 million registered members who upload more than 3.5 million new images daily[2]. We collected a sample of the follower network composed by the nearly $40M$ public Flickr profiles that are opted-in for research studies and by the $500M+$ links that connect them. Links carry the timestamp of their creation and they span approximately 12 years ending March 2016. Similar to Tumblr, Flickr has a contact recommendation module. However, we do not have access to recommendation data and we cannot measure their effect on the link creation process.
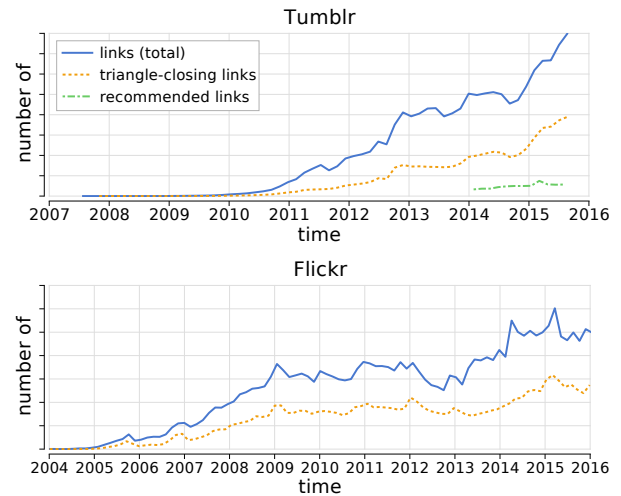
Figure 3: Number of links, triangle-closing links, and recommended links (Tumblr only) created each day, by the nodes in our sample, during the whole lifespan of the platforms.

3.3 Concepts and notation

Graph and ego-network. Consider a follower graph $\mathcal{G}$ composed by a set of nodes $\mathcal{N}$ and a set of directed edges[3] $\mathcal{E} \in \mathcal{N} \times \mathcal{N}$. When building the follower graph, we draw an edge from node i to node j if i has followed j at any time. The *ego-network* $\mathcal{G}_i$ of node $i \in \mathcal{N}$ is the subgraph induced by i's out-neighbors $\Gamma_{out}(i)$ [22]. Formally: $\mathcal{G}_i = (\mathcal{N}_i, \mathcal{E}_i)$, where $\mathcal{N}_i = \Gamma_{out}(i)$, and $\mathcal{E}_i = \{(j,l) \in \mathcal{E} |\ j \in \Gamma_{out}(i) \land l \in \Gamma_{out}(i)\}$. Note that the ego-network does *not* include the links between the ego i and its neighbors.

Structural graph metrics in time. The temporal trace of link creations allows us to build a time graph [32] and to recover the structural properties of nodes and links at any point in time. The superscript t applied to any indicator means that the metric refers to a snapshot of the graph at time t. For example, the neighbor set of node i at time t is denoted as $\Gamma^t(i)$ and its degree as $k^t(i)$. When studying the evolution of ego-networks in isolation, we will consider time on a discrete scale where each event corresponds to the n^{th} node being added to the ego-network. We will use the letter n to denote time passing on this discrete scale. All the graphs we consider are directed, so we use the definition of triangle closure adapted to directed graphs [43]: a new link created between i and j at time t closes a directed triangle if $\exists l \in \mathcal{N} | l \in \Gamma^t_{out}(i) \land l \in \Gamma^t_{in}(j)$.

Spontaneous vs. recommended links. We distinguish links that are created for effect of a recommendation from those that are not. We call the links in the first group *recommended* and the ones in latter *spontaneous*.

3.4 Data overview

The (in/out)degree distributions together with their average values (μ) are shown in Figure 2. As expected, all distributions are broad, with values spanning several orders of magnitude. The out-degree distribution in Tumblr is capped
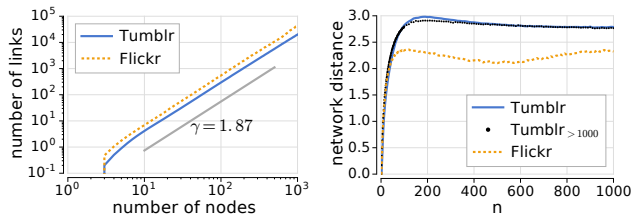
Figure 4: Left: average number of ego-network links vs. number of nodes; best fitting power-law exponent is reported as reference. Right: average network distance after the n^{th} node is added; the black dotted line is obtained considering only Tumblr ego-networks with at least 1000 nodes.
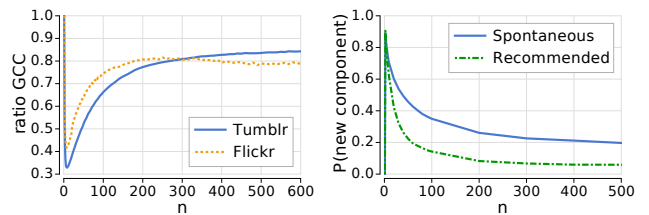

Figure 5: Left: average ratio of nodes in the giant weakly connected component (GCC) as new nodes are added to the ego network. Right: probability that the n^{th} node included in the ego-network spawns a new disconnected component, computed for spontaneous and recommended nodes.

at 5000 because the platform imposes an upper bound on the number of blogs a user can follow. On average, Tumblr users are more connected than Flickr users, with average in- and out-degree 5 and 3 times larger than Flickr, respectively.

Figure 3 plots the number of links created over the course of the platforms' life. Both networks have experienced a noticeable growth. For Tumblr, we can plot the time series of recommended link creations occurred after January 2015. We also calculate the set of links that close at least one triangle. In Tumblr, the first sharp increase in the number of triangle-closing links is found between 2012 and 2013. That is determined by the introduction of a new link recommender system. A similar pattern has been observed in Facebook after the introduction of the "people you may know" module [60]. About 27% of recommended links do not close any triangle: those are recommendations based on the user's topical profile only.

4. EVOLUTION OF EGO-NETWORKS

4.1 Diameter and connected components

The growth of social networks is associated with three changes in their macroscopic structure: densification, diameter shrinking, and inclusion of almost all nodes in a single giant connected component [34]. It is unknown whether the same properties hold at ego-network level.

Q1: How do density, diameter, and component structure evolve as the ego-network grows?

Like global networks, ego-networks become denser in time. Ego-networks obey a densification power law, for which the number of links scales superlinearly with the number of nodes $|\mathcal{E}_i| \sim |\mathcal{N}_i|^\gamma$ (Figure 4, left). The exponent that best defines the scaling in both platforms is $\gamma = 1.87$.

More surprisingly, densification does not always lead to the emergence of a single giant connected component covering the whole graph. On average[4], the largest component's size relative to the network size grows as new nodes join (Figure 5, left), but it stabilizes around 0.8 for networks of 200 nodes or more. The number of components grows sublinearly with the number of nodes (not shown). More notably, the diameter shows little signs of shrinking. The *network distance*, computed as the average distance between all pairs of nodes,[5] experiences a three-phases evolution (Fig-

---

[4]Results are qualitatively similar when considering the median.

[5]Computed on an undirected version on the graph. Similar results are obtained using diameter or effective diameter.

ure 4, right). First, at the beginning of the ego-network life, it expands rapidly (*exploration*); then, it starts shrinking slightly (*consolidation*) before asintotically converging to a stable value (*stabilization*). This trend is very different from the sharp diameter decline that characterizes social graphs. We speculate that the consolidation phase might be connected with the intrinsic human limitation to maintain large social groups, as theorized by Robin Dunbar [20]. When the ego's social neighborhood exceeds the size that is cognitively manageable by a person (roughly, 150 to 200 individuals), a compensation effect might be triggered: new contacts are not anymore sought further away from the social circles that have been already established, putting an end to the exploration phase. This happens at $n = 140$ in Flickr and $n = 190$ in Tumblr, values that are compatible with Dunbar's theory.

The addition of recommended nodes has a different effect on the ego-network expansion, compared to spontaneous ones. New recommended contacts tend to be closer to existing ego-network members. At fixed network size, the addition of new recommended nodes increases the network distance 5% less, on average, than a spontaneous node addition. Ego-networks with at least a recommended node have smaller network distance than ego-networks that grew fully spontaneously; this difference varies with the size, being only 2% smaller for networks under 50 nodes up to 10% smaller for networks with 200 nodes or more. Also, spontaneous nodes have far higher chances, compared to recommended ones, to spawn a new component disconnected from the rest of the ego-network (Figure 5, right).

Accounting for amalgamation effects. To explore evolutionary trends of ego-networks, we rely on aggregate analysis: an indicator is measured on an ego-network when its n^{th} node is added and then it is averaged across all ego-networks. Trends are discovered as n grows (e.g., diameter in Figure 4, right). This approach may yield misleading results because averages computed at different values of n are obtained from different sample sets. This problem is known as the *Simpson's Paradox* [46] and it is usually addressed by fixing the sample set [12]. To account for it, every time we perform an evolutionary analysis as n varies in $[1, n_{max}]$, we compare results obtained in two settings: the first using the full dataset and the latter considering only the subset of ego-networks that reached at least size n_{max}. The results are only slightly different across the two settings, for all the indicators analyzed. For the sake of brevity, we report just one example of such comparison. In Figure 4 right, the di-
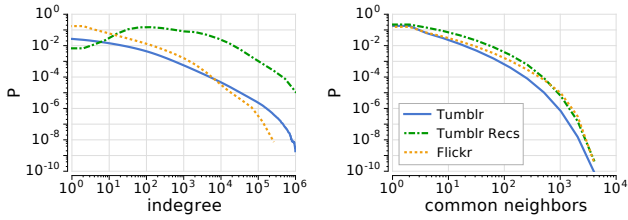
Figure 6: Distribution of nodes' popularity (indegree) and similarity with ego (common neighbors) at the time of their inclusion in the ego-network.
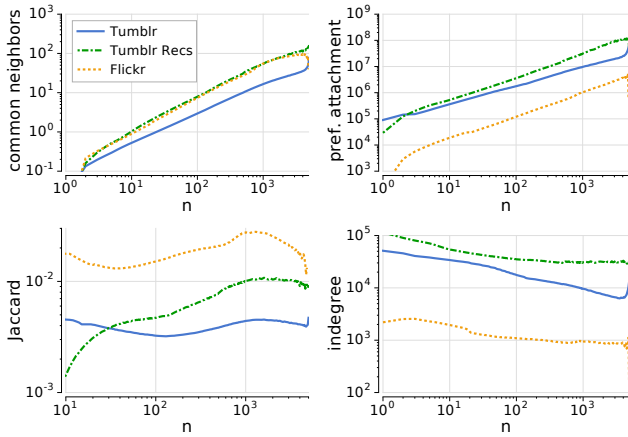

Figure 7: Average similarity and popularity indicators of n^{th} node added to the ego-network: common neghbors between the ego and the newly added nodes, Jaccard similarity between their neighbors sets, preferential attachment indicator ($k_{out}(ego) \cdot k_{in}(newnode)$), and indegree of the new node.

ameter evolution for Tumblr ego-networks that reached at least size 1000 is very similar to the trend found when all ego-networks are considered.

4.2 Popularity vs. similarity

The process of link creation in online social networks is driven by two main factors: popularity (that leads to preferential attachment [38]) and similarity (that leads to homophily [2]). At network scale, their relative weight in predicting the creation of new links might vary depending on the type of social network [28]. At microscopic scale, it is still unclear how popularity and similarity impact the selection of new nodes in ego-networks, and how their relative importance varies in time.

Q2: How do the criteria of neighbor selection change as the ego-network grows?

We select two simple (yet widely-used) proxies of popularity and similarity. Given an ego i who has added j as its neighbor at time t, we consider the alter's indegree $k_{in}^t(j)$ as an indicator of its popularity and the number of common neighbors between the ego i and the alter j, $CN^t(i,j) = |\Gamma_{out}^t(i) \cap \Gamma_{in}^t(j)|$, as a measure of similarity. Drawing the distributions of k_{in}^t and $CN^t(i,j)$ (Figure 6), we observe that the range of values is very broad in both platforms. The CN distributions suffer from cut-offs (around $CN = 200$) caused by the scarcity of nodes with hundreds of common
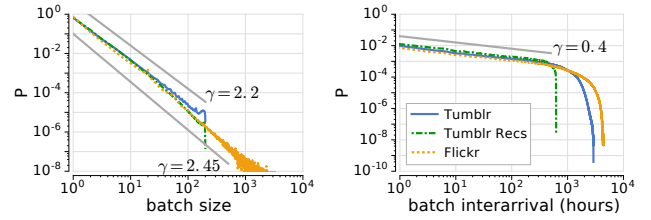

Figure 8: Distributions of batch size (s_b) and batch interarrival time (τ_b). Best fitting power law exponents reported as reference.
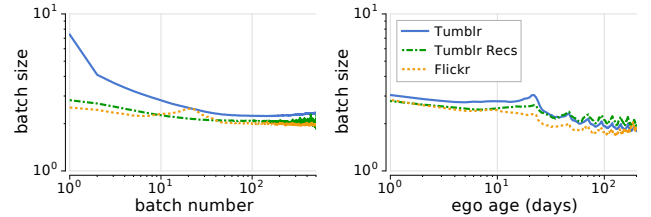

Figure 9: Batch size as a function of time, measured as the number of batches created or as the ego's age measured in number of days. The Tumblr Recs curves summarize the trend for batches containing at least one recommended node.

neighbors or more. Recommended Tumblr nodes yield distributions that are skewed towards higher values because the recommender picks by design those profiles that are popular and well-connected to the ego's neighbors.

As new nodes are added to the ego-network, the number of their common connections with the ego naturally increases (and so does the preferential attachment indicator, as expected). The Jaccard similarity between their neighbor sets oscillates, increasing when the ego-network's size is in the interval $[100, 1000]$ and decreasing otherwise. The popularity of new ego-network members, computed as their indegree in the social network, decreases as the ego-network grows. A summary of all the indicators is given in Figure 7.

All the trends are similar, yet shifted towards higher values, when considering recommended nodes only. In short, recommended nodes tend to share more contacts with the ego and to be more popular, which corroborates previous observations about link recommendations being beneficial mostly to popular nodes [48]. The indicator that differs the most is the Jaccard similarity, that increases monotonically with n for recommended contacts.

4.3 Temporal activity

Creation of links is not uniform in time. Previous literature found evidence that, globally, the creation of links happens in bursts [32, 30]. At a local level, we aim to learn how often and in which phases of the ego-network life users select new neighbors.

Q3: When do ego-networks expand?

To measure how much node additions to an ego-network are concentrated in short periods of time, we resort to basic session analysis to group together temporally-contiguous events. As is standard practice in the analysis of browsing behaviour [47], we split user sessions by timeout: a session starts when a new node is added to the ego-network and
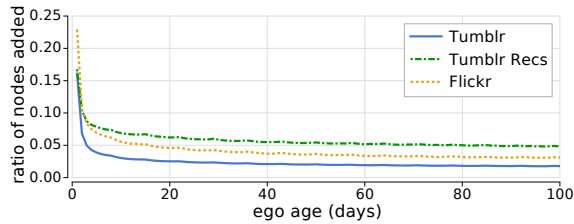
Figure 10: Average portion of links created in the first 100 days of the ego's life, relative to the final ego-network size. Only nodes who have created links for at least 6 months are considered.

Figure 11: Left: distribution of number of communities in ego-networks. Right: average size of communities as they appear in the ego-network; 95% confidence intervals are shown.

ends when no other node has been added for 25 minutes. We call *batch* a set of nodes added in a single session.

We compute the average *batch size* s_b and the session *interarrival time* τ_b, namely the time (hours) elapsed from the session's end to the next session's start. The process of batch creation is bursty when *i*) there are strong temporal heterogeneities in the interarrival time, and *ii*) consecutive link creations are not independent events. A standard practice to assess those conditions is to measure the decay of the probability density functions for s_b and τ_b: power law decays in the form $P(x) \sim x^{-\gamma}$ indicate burstiness [29]. The distribution of batch size s_b follows a power-law trend, with exponents 2.2 and 2.45 in Tumblr and Flickr, respectively (Figure 8 left). In Flickr, the size scales freely as there are no boundaries preventing the addition of any number of contact. In Tumblr, we observe a sharp cutoff at 200 as the service policy enforces a maximum limit of 200 link creations per user per day. The decay of τ_b is similar on both platforms, with initial intervals fitting a power-law with exponent $\gamma = 0.4$, followed by exponential cutoffs due to the finite time window (Figure 8, right).

If we consider only sequences of batches containing at least one recommended link, we see that recommendations are associated with the creation of less links per session, but at higher rate. The average batch size is 2.18 in Flickr and 2.67 in Tumblr; Tumblr batches with recommended links are 12% smaller (average size 2.35). The median interarrival time is relatively high in both platforms —12 days in Tumblr, 2 weeks in Flickr— but only 5 days for pairs of consecutive batches containing recommended links. No causal claim connecting recommendations and rate of link creation can be made, as a number of confounding factors could influence this trend (e.g., users who are more active might more naturally engage in recommendations). However, this result provides partial evidence that recommendations might contribute to alter the natural time scale of link creation.

The average batch size decreases as time passes and the ego-network grows (Figure 9). After the first 30 days (or the first 20-30 batches created), the batch size stabilizes around 2. A similar decreasing trend is also found for the interarrival time τ_b (not shown). This suggest that nodes tend to build most of their ego-network in the first stages of their life. To confirm that, we compute the average daily ratio of the total number of the ego-network's nodes added to in the first 100 days of the ego's life. We only consider users whose link creation activity spans at least 6 months, to avoid biases introduced by users with short lifespan. As expected, a big chunk of nodes are typically added in the first days of activity (Figure 10). This finding adds nuance to previous wo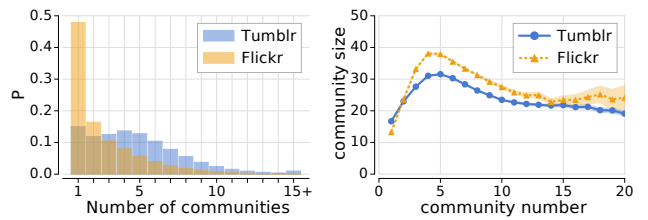rk on temporal graphs. Studies on Flickr using a coarser temporal granularity found that the raw number of new links created by the ego in time is uniform over time [33]. Here we find that the uniform trend starts only after an initial spike of link creations.

4.4 Community formation

Ego-networks have a clear community structure because people tend to interact with multiple social circles (e.g., school friends, family members) that are typically weakly connected to one another [37]. We ask about the role of these communities in the graph evolution.

Q4: Is the ego-network growth driven by the boundaries of its communities?

The ego-network may follow a depth-first expansion pattern with respect to communties, in which the ego preferentially connects to nodes belonging to a community before exploring others. In Flickr, for example, a person could first follow all the accounts of family members and then those of a photography club. Alternatively, the ego-network may expand either breadth-first, picking new nodes in a round-robin fashion, or regardless of the community structure. In social network analysis research we know little evidence in support of any of these scenarios. In the context of web navigation and search, in-depth exploration of content is often most effective and cognitively more natural [18, 50]; we hypothesize that the same holds for community exploration.

Measuring the extent to which any of these scenarios reflects people's behaviour is challenging, as in reality communities can be overlapping and change their boundaries as the graph grows. Leaving more advanced measurements for future work, we assume a static, hard partitioning of nodes in communities. For all ego-networks, we compute non-overlapping communities[6] at time $t = T_{end}$ (the most recent snapshot in our data), heuristically filtering out small ego-networks with less than 5 members. Ego-networks are often composed by few communities, rarely more than 10 (Figure 11, left).

To assess to what extent communities emerge over time in an orderly fashion, we rank them by the time the ego has created connections with their nodes. Specifically, we first sort all the nodes by the time they are added to the ego-network (e.g., $[j_1, j_2, j_3, j_4]$). We then replace nodes with the communities they belong to (e.g., $[c_1, c_1, c_2, c_1]$), rank communities by the median position p of their occurrences in that vector (e.g., $p(c_1) = 2, p(c_2) = 3$; c_1 is ranked first, c_2 is ranked last), and finally replacing the communities with their respective ranks, thus obtaining a

---

[6]We use the community detection algorithm by Waltman and Eck [52] that is an optimization of the Louvain method [14].
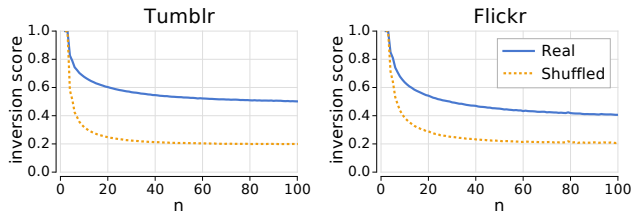
Figure 12: Average inversion score of ego-network communities as new nodes are added, compared to a randomized null-model.
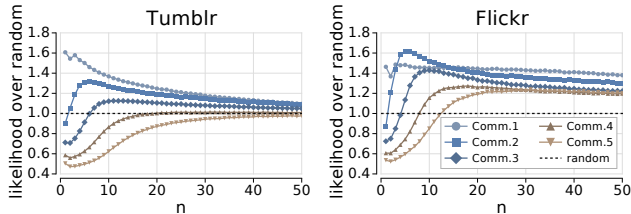

Figure 13: Average likelihood (over random chance) that the n^{th} node in a ego-network belongs to the k^{th} community.

sequence of community ranks $\mathcal{R}$ (e.g., [1,1,2,1]). The intuition is that a community half of whose members has been added to the ego-network by time t comes temporally before any other whose majority of nodes are still outside the ego-network at the same time t. If exploration of communities is purely in-depth, $\mathcal{R}$ is fully sorted. The sortedness of a list $L = \{x_1, ..., x_m\}$ can be measured by its *inversion score*:

$$inv(L) = 1 - \frac{2 \cdot |\{(x_i, x_j)|i < j \wedge x_i > x_j\}|}{\binom{|L|}{2}} \in [-1, 1]$$

$inv = 1$ indicates sortedness, $inv = -1$ inverse ordering, and $inv = 0$ randomness. In Figure 12 we plot the average inversion score of $\mathcal{R}$ against the ego-network size. To account for the community size heterogeneity, we compare it with a null-model where the elements in $\mathcal{R}$ are randomly reshuffled. The inversion score quickly stabilizes as the network grows and it has values that are consistently higher (double or more) than the null-model's, supporting the hypothesis that communities tend to be explored in depth, one after the other.

Further evidence can be provided by measuring the probability that the n^{th} node added to the ego-network belongs to the k^{th} community in the ranking, normalized by the probability in the null-model; values higher than 1 indicate above-chance likelihood of a node being in a given community. Figure 13 shows the average normalized likelihood of the first 50 nodes to belong to the first 5 communities in the rank. Curves for increasing values of k emerge above the randomness threshold one after the other, which backs the hypothesis of communities being explored in-depth.

The size of a community (measured at time T_{end}) varies with its temporal rank (Figure 11, right). On average, people create increasingly larger communities up to the fifth one; from the sixth one on, new communities added become smaller and smaller.

4.5 Recommendations diversity

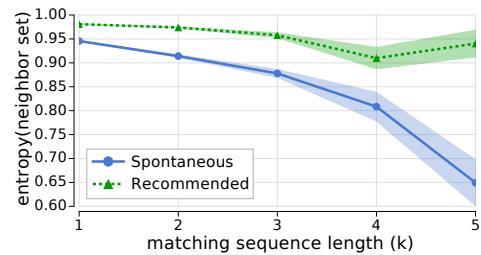Our analysis shows that the statistical properties of recommended links are different from those of spontaneous ones.

Figure 14: Matching experiment. Average entropy of the neighbor sets at step $k + 1$ for groups sharing a matching sequence of length k. 95% confidence intervals are shown.

It is known that recommender systems postively affect user engagement, in terms of time spent, content consumption, and user contribution [23]. In agreement with established knowledge, we have found that users exposed to recommendations create more links, more frequently. It is harder to assess whether recommendations foster or limit access to diverse types of content. The academic debate about recommendations being the bane or boon of social media is still very lively [42, 10, 44], with evidence brought in support of the two views. We aim to provide further evidence to shed light on this point in the context of link recommenders.

Q5: Do link recommendations foster diversity?

It is hard to infer causality from observational data. Matching is a statistical technique that is used to evaluate the effect of a treatment on a dependent variable by comparing individuals who have received the treatment with others with similar observable features who did not receive it. The more similar the paired individuals and the higher the number of pairs, the higher the confidence of estimating the cause of the treatment on the dependent variable.

We conduct a matching experiment to measure if people who follow recommendations end up having ego-networks more similar to one another than if they were to ignore recommendations. We arrange people in *matching groups* containing users who are nearly identical in terms of their local connectivity. We assign to the same matching group users who i) registered to Tumblr less than 30 days apart, and ii) whose very first k neighbors are the same and have been added in the same order to their ego-networks. Each matching group is then split in two subgroups: a *treatment* group of users whose $k+1^{st}$ contact is a recommended one, and a *control* group of all the remaining users, whose $k+1^{st}$ contact has been created spontaneously. The variety of contacts created at step $k + 1$ is our dependent variable. If the variety measured in one group is significantly different from the other, we can attribute the divergence to the effect of the recommendation, as the initial conditions of the two groups are virtually identical. For example, if we found that the variety of nodes in the treatment group is lower than the one measured on the control group, we would conclude that recommendations conform the process of link creation by inducing users to follow a more restricted set of accounts compared to what would happen by spontaneous user behavior. We measure diversity of contacts through their entropy. To ensure a fair comparison that accounts for size heterogeneity, we use normalized entropy $\widehat{H}$. Given a bag of nodes X of size N, where $p(x)$ is the number of occurrences of node

Model	AUC	F-Score
Baseline	0.893	0.813
+ age	0.938	0.864
+ k_{out}	0.897	0.817
All	**0.943**	**0.87**

Table 1: Link prediction results.

Figure 15: Feature importance.

Figure 16: Distribution of structural features of spontaneous and recommended links. Given directed links in the form (i, j) we show the boxplots of the distributions of $k_{in}(j)$, $PA(i, j)$, $CN(i, j)$, and Jaccard(i, j).
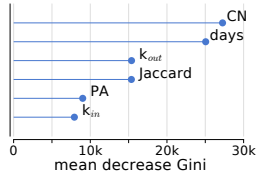
$x \in X$ divided by N, the normalized entropy is defined as:

$$\widehat{H}(X) = \sum_{x \in X} \frac{p(x) \cdot log_2(p(x))}{log_2(N)} \in [0, 1].$$

Figure 14 shows the results for a total of approximately 25K matching groups, for $k \in [1, 5]$; requiring k identical links is a too strong requirement for larger k. Every point is the average of all the matching groups for a given k. First, we observe that, the higher the k, the lower the entropy. That is expected: the higher the number of common neighbors, the more likely the next selected neighbor will be the same. Last, most importantly, the variety of the treatment group is always higher; this indicates that recommendations foster diversity. Although it is difficult to pin down the exact reasons why this happens, we provide a possible interpretation. Even if the list of recommended contacts was the same for all the users in the treatment group, the reshuffling of the top recommended contacts that Tumblr implements in the link recommender widget introduces asymmetries across users. More generally, we could hypothesize that the recommender system exposes users to a wider set of potential contacts than the ones they would be exposed to by browsing or searching on the site, thus providing a wider spectrum of options and, in turn, a more diverse set of individual choices.

To test the robustness of the results, we explored some possible alternatives in the setup of the matching experiment. Specifically we: *i*) measured the entropy at $k + 2$ instead of $k + 1$ (we leave the $k + n$ generalization for future work); *ii*) selected only control and treatment groups with at least $m \in [2, 15]$ members (the results reported are for $m = 5$); *iii*) randomly downsampled the larger group to match the size of the smaller one, to balance the size of the two; *iv*) run two independent experiments including in the treatment group only users whose recommended contact had 1) at least one common neighbor (i.e., the recommendation is provided based on network topology features) or 2) no common neighbors (i.e., the recommendation is provided on a topical basis). The absolute values vary slightly across setups, but the qualitative results remain the same.

5. IMPACT ON LINK PREDICTION

Our analytical results have direct implications on how link recommender systems can be enhanced to provide more effective suggestions. Next, we discuss two prediction experiments that aim to answer two research questions.

Q6: To what extent temporal features improve the ability to predict new links?

The previous analysis showed that egos add new neighbors to their network with criteria that change in time. The resulting hypothesis is that link recommendations that adapt to the current evolutionary stage of the ego-network could
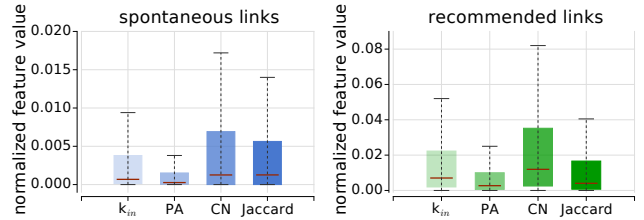
gain effectiveness. To test such hypothesis, we run a prediction experiment.

We consider a snapshot of the Tumblr social network at an arbitrary time t (January 1^{st} 2015). To build a training set, we sample 200K node pairs (i, j) that are not directly connected but with at least one directed common neighbor (i.e., there is a directed path of length 2 from i to j). Half of the pairs will be directly connected by a link from i to j before T_{end} (positive examples), the remaining half will remain disconnected (negative examples). For every pair, we extract six simple features: i's outdegree ($k_{out}(i)$), j's indegree ($k_{in}(j)$), preferential attachment (PA $= k_{out}(i) \cdot k_{in}(j)$), common neighbors (CN $= |\Gamma_{out}^t(i) \cap \Gamma_{in}^t(j)|$), Jaccard similarity between neighbor sets (Jaccard $= \frac{|\Gamma_{out}^t(i) \cap \Gamma_{in}^t(j)|}{|\Gamma_{out}^t(i) \cup \Gamma_{in}^t(j)|}$), and i's age measured in number of days elapsed from i's profile creation to t. Age and i's outdegree are the two temporal features whose effectiveness we want to investigate: one measures time on a continuous scale, the other on the discrete scale of link creation events.

The simple features above can predict the node pair class very accurately, and a summary of this evaluation (10-fold cross validation using random forest) is given in Table 1. The model trained on the full set of features yields an AUC of 0.943 and a F-measure of 0.87, which is an improvement of 5.6%/7% in terms of AUC/F-measure over a baseline model that does not consider temporal features. Adding the feature k_{out} to the baseline model yields only a slight improvement in accuracy (0.45%/0.49%), while considering age improves the accuracy in a more consistent way (5.04%/6.27%). Figure 15 summarizes the importance of the features in this link prediction setting by measuring the *mean decrease Gini*, the average gain of purity achieved when splitting on a given variable (the higher, the better) [36]. In line with previous work [56], this analysis further confirms the importance of the temporal features and found that time matters when recommending new contacts.

Q7: Is it possible to limit the bias of the recommender system while keeping its high accuracy?

In Section 4.2 we observed that the distribution of nodes' popularity and of their structural similarity with ego are statistically different when considering recommended vs. spontaneous links. To further investigate this difference, we analyze the distribution of all structural features over a sample (200k) of recommended and spontaneous links, equally represented. For sake of presentation, we normalize the value of each feature on a scale 0 to 1, by dividing the raw feature value by its maximum value in the whole 200k sample. As shown in Figure 16, spontaneous links tend to exhibit less de-

gree of structural overlap than recommended links (the median Jaccard value on recommended links is 4 times larger than the value recorded on spontaneous ones). The same observation holds when analyzing the node popularity; the median in-degree of the target node on recommended links is one order of magnitude higher than the corresponding value computed on spontaneous links.

Learning to what extent it is possible to automatically tell recommended links and spontaneous links apart would allow us to train new recommender systems to suggest links whose properties better adhere to the natural criteria that people follow when adding new contacts. To gauge this possibility, we run a second prediction experiment with the same features and setup of the previous one but with a different selection of positive and negative examples.

We pick $100K$ pairs that will be connected in the future through a spontaneous link as positive examples, and as many pairs that will be connected through a recommended link as negative examples. A random forest classifier is able to distinguish the two classes pretty accurately (AUC = 0.795, F-measure = 0.721).

In a more realistic scenario, the recommender should learn to recognize the space between recommended links and truly negative examples (links that are never formed). To model that, we add to the training set $100K$ negative pairs that will not be connected at any time in the future. This setting achieves a better performance (AUC = 0.823, F-measure = 0.771) with around 90% accuracy on the negative class and 55% on the positive one. In short, even if very basic structural and temporal features are used, it is possible to effectively use the output of current link recommenders to train new recommenders that smooth the algorithmic bias and produce suggestions that better simulate the spontaneous process of link selection.

6. CONCLUSIONS

We have provided a large-scale analysis of ego-network evolution on two online platforms, exposing the dynamics of their bursty evolution, community-driven growth, diameter expansion, and selection of new nodes based on a time-varying interplay between similarity and popularity. By studying the set of Tumblr links created as a result of algorithmic suggestions, we find that recommended links have different statistical properties than spontaneously-generated ones. We also find evidence that link recommendations foster network diversity by leading nodes that are structurally similar to choose different sets of new neighbors.

Our work has some limitations. Flickr and Tumblr are mainly interest networks, where people follow each other based on topical tastes. Some of the results we report here might not generalize to social networks that aim mostly at connecting people who know each other in real life (e.g., Facebook, LinkedIn). Also, we only consider the network structure and disregard any notion of node profile, including posting activity in time and user-generated content; that information would help to further detail the dynamics of ego-network expansion with respect to other dimensions such as topical similarity between profiles.

Our results have a number of practical implications. We provide further evidence that in online social networks not all links are created equal; network analysts who produce network growth models based on the observation of online social networks' longitudinal traces should consider weight-

ing links that emerge spontaneously different from those that are created algorithmically. Through our prediction experiments, we also provide a hint about how link recommender systems could incorporate signals on the ego-network's evolutionary stage to improve the quality of suggestions. We hope our work provides yet another step towards a better understanding of the evolutionary dynamics of social networks.

7. ACKNOWLEDGMENTS

8. REFERENCES

[1] Y. Y. Ahn, S. Han, H. Kwak, S. Moon, and H. Jeong. Analysis of topological characteristics of huge online social networking services. In *WWW*, 2007.

[2] L. M. Aiello, A. Barrat, R. Schifanella, C. Cattuto, B. Markines, and F. Menczer. Friendship prediction and homophily in social media. *ACM TWEB*, 2012.

[3] V. Arnaboldi, M. Conti, M. La Gala, A. Passarella, and F. Pezzoni. Ego network structure in online social networks and its impact on information diffusion. *Comp. Comm.*, 2016.

[4] V. Arnaboldi, M. Conti, A. Passarella, and F. Pezzoni. Analysis of ego network structure in online social networks. In *SocialCom*, 2012.

[5] V. Arnaboldi, M. Conti, A. Passarella, and F. Pezzoni. Ego networks in twitter: an experimental analysis. In *InfoCom*, 2013.

[6] V. Arnaboldi, R. Dunbar, A. Passarella, and M. Conti. Analysis of co-authorship ego networks. In *Network Science*. Springer, 2016.

[7] L. Backstrom, D. Huttenlocher, J. Kleinberg, and X. Lan. Group formation in large social networks: membership, growth, and evolution. In *KDD*. ACM, 2006.

[8] L. Backstrom and J. Leskovec. Supervised random walks: Predicting and recommending links in social networks. In *WSDM*, 2011.

[9] B. Bahmani, A. Chowdhury, and A. Goel. Fast incremental and personalized pagerank. *Proc. VLDB Endow.*, 2010.

[10] E. Bakshy, S. Messing, and L. A. Adamic. Exposure to ideologically diverse news and opinion on facebook. *Science*, 2015.

[11] A.-L. Barabâsi, H. Jeong, Z. Néda, E. Ravasz, A. Schubert, and T. Vicsek. Evolution of the social network of scientific collaborations. *Physica A: Statistical mechanics and its applications*, 311(3), 2002.

[12] S. Barbosa, D. Cosley, A. Sharma, and R. M. Cesar Jr. Averaging gone wrong: Using time-aware analyses to better understand behavior. In *WWW*. ACM, 2016.

[13] A. Biswas and B. Biswas. Investigating community structure in perspective of ego network. *Expert Syst. Appl.*, 2015.

[14] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre. Fast unfolding of communities in large networks. *Journ. of Stat. Mech.*, 2008.

[15] Y. Chang, L. Tang, Y. Inagaki, and Y. Liu. What is tumblr: A statistical overview and comparison. *SIGKDD Explor. Newsl.*, 2014.

[16] J. Chen, W. Geyer, C. Dugan, M. Muller, and I. Guy. Make new friends, but keep the old: recommending people on social networking sites. 2009.

[17] E. M. Daly, W. Geyer, and D. R. Millen. The network effects of recommending social connections. In *RecSys*, 2010.

[18] P. De Bra, G.-J. Houben, Y. Kornatzky, and R. Post. Information retrieval in distributed hypertexts. In *RIAO*, 1994.

[19] A. De Salve, M. Dondio, B. Guidi, and L. Ricci. The impact of user's availability on on-line ego networks: a facebook analysis. *Comp. Comm.*, 2016.

[20] R. Dunbar. *Grooming, gossip, and the evolution of language*. Harvard University Press, 1998.

[21] S. L. Feld. Why your friends have more friends than you do. *Am. Jour. of Soc.*, 1991.

[22] L. C. Freeman. Centered graphs and the structure of ego networks. *Math. Soc. Sci.*, 1982.

[23] J. Freyne, M. Jacovi, I. Guy, and W. Geyer. Increasing engagement through early recommender intervention. In *RecSys*, 2009.

[24] L. Garton, C. Haythornthwaite, and B. Wellman. Studying online social networks. *Journ. of Com. Med. Comm.*, 1997.

[25] M. Grbovic, V. Radosavljevic, N. Djuric, N. Bhamidipati, and A. Nagarajan. Gender and interest targeting for sponsored post advertising at tumblr. In *KDD*, 2015.

[26] P. Gupta, A. Goel, J. Lin, A. Sharma, D. Wang, and R. Zadeh. Wtf: The who to follow service at twitter. In *WWW*, 2013.

[27] S. Gupta, X. Yan, and K. Lerman. Structural properties of ego networks. In *SBP-BRiMS*. Springer, 2015.

[28] M. A. Hasan and M. J. Zaki. *Social Network Data Analytics*, chapter A Survey of Link Prediction in Social Networks. Springer US, Boston, MA, 2011.

[29] M. Karsai, K. Kaski, A.-L. Barabási, and J. Kertész. Universal features of correlated bursty behaviour. *Scientific reports*, 2012.

[30] R. Kikas, M. Dumas, and M. Karsai. Bursty egocentric network evolution in skype. *Soc. Net. An. and Min.*, 2013.

[31] G. Kossinets and D. J. Watts. Empirical analysis of an evolving social network. *Science*, 2006.

[32] R. Kumar, J. Novak, P. Raghavan, and A. Tomkins. On the bursty evolution of blogspace. In *WWW*, 2003.

[33] J. Leskovec, L. Backstrom, R. Kumar, and A. Tomkins. Microscopic evolution of social networks. In *KDD*, 2008.

[34] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graphs over time: Densification laws, shrinking diameters and possible explanations. In *KDD*, 2005.

[35] D. Liben-Nowell and J. Kleinberg. The link prediction problem for social networks. In *CIKM*.

[36] G. Louppe, L. Wehenkel, A. Sutera, and P. Geurts. Understanding variable importances in forests of randomized trees. In *NIPS*. Curran Associates Inc., 2013.

[37] J. Mcauley and J. Leskovec. Discovering social circles in ego networks. *TKDD*, 2014.

[38] A. Mislove, H. S. Koppula, K. P. Gummadi, P. Druschel, and B. Bhattacharjee. Growth of the flickr social network. In *WOSN*, 2008.

[39] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee. Measurement and analysis of online social networks. In *IMC*, 2007.

[40] S. A. Muhammad and K. Van Laerhoven. Duke: A solution for discovering neighborhood patterns in ego networks. In

[41] G. Palla, A.-L. Barabási, and T. Vicsek. Quantifying social group evolution. *Nature*, 446(7136), 2007.

[42] E. Pariser. *The Filter Bubble: How the New Personalized Web is Changing what We Read and how We Think*. Penguin Books, 2012.

[43] D. M. Romero and J. M. Kleinberg. The directed closure process in hybrid social-information networks, with an analysis of link formation on twitter. In *ICWSM*, 2010.

[44] A. Sharma, J. M. Hofman, and D. J. Watts. Estimating the causal impact of recommendation systems from observational data. In *EC*, 2015.

[45] D. Shin, S. Cetintas, K.-C. Lee, and I. S. Dhillon. Tumblr blog recommendation with boosted inductive matrix completion. In *CIKM*, 2015.

[46] E. H. Simpson. The interpretation of interaction in contingency tables. *Journal of the Royal Statistical Society. Series B*, 1951.

[47] A. Spink, M. Park, B. J. Jansen, and J. Pedersen. Multitasking during web search sessions. *Inf. Proc. and Man.*, 2006.

[48] J. Su, A. Sharma, and S. Goel. The effect of recommendations on network structure. In *WWW*, 2016.

[49] C. Tan and L. Lee. All who wander: On the prevalence and characteristics of multi-community engagement. In *WWW*. ACM, 2015.

[50] L. Tauscher and S. Greenberg. How people revisit web pages. *Int. J. Hum. Comput. Stud.*, 1997.

[51] J. Ugander, B. Karrer, L. Backstrom, and C. Marlow. The anatomy of the facebook social graph. *CoRR*, abs/1111.4503, 2011.

[52] L. Waltman and N. J. Eck. A smart local moving algorithm for large-scale modularity-based community detection. *Eur. Phys. J. B*, 2013.

[53] L. Weng and T. M. Lento. Topic-based clusters in egocentric networks on facebook. In *ICWSM*, 2014.

[54] C. Wilson, B. Boe, A. Sala, K. P. Puttaswamy, and B. Y. Zhao. User interactions in social networks and their implications. In *EuroSys*, 2009.

[55] J. Yang and J. Leskovec. Patterns of temporal variation in online media. In *WSDM*, 2011.

[56] Y. Yang, N. Chawla, Y. Sun, and J. Hani. Predicting links in multi-relational and heterogeneous networks. In *ICDM*. IEEE, 2012.

[57] D. Yin, L. Hong, X. Xiong, and B. D. Davison. Link formation analysis in microblogs. In *SIGIR*, 2011.

[58] G. Yuan, P. K. Murukannaiah, Z. Zhang, and M. P. Singh. Exploiting sentiment homophily for link prediction. In *RecSys*, 2014.

[59] W.-X. Zhou, D. Sornette, R. A. Hill, and R. I. Dunbar. Discrete hierarchical organization of social group sizes. *Royal Society of London B*, 2005.

[60] M. Zignani, S. Gaito, G. P. Rossi, X. Zhao, H. Zheng, and B. Y. Zhao. Link and triadic closure delay: Temporal metrics for social network dynamics. In *ICWSM*, 2014.