

People Are Strange When You're a Stranger: Impact and Influence of Bots on Social Networks

Luca Maria Aiello, Martina Deplano, Rossano Schifanella, Giancarlo Ruffo

Computer Science Department
Università degli Studi di Torino
Torino, Italy

Abstract

Bots are, for many Web and social media users, the source of many dangerous attacks or the carrier of unwanted messages, such as spam. Nevertheless, crawlers and software agents are a precious tool for analysts, and they are continuously executed to collect data or to test distributed applications. However, no one knows which is the real potential of a bot whose purpose is to control a community, to manipulate consensus, or to influence user behavior. It is commonly believed that the better an agent simulates human behavior in a social network, the more it can succeed to generate an impact in that community. We contribute to shed light on this issue through an online social experiment aimed to study to what extent a bot with no trust, no profile, and no aims to reproduce human behavior, can become popular and influential in a social media. Results show that a basic social probing activity can be used to acquire social relevance on the network and that the so-acquired popularity can be effectively leveraged to drive users in their social connectivity choices. We also register that our bot activity unveiled hidden social polarization patterns in the community and triggered an emotional response of individuals that brings to light subtle privacy hazards perceived by the user base.

1 Introduction

The capillary diffusion of online social media and the growing number of mobile devices connected to the Internet have determined an unprecedented expansion of the pervasiveness of the Web in real life. Online and offline spheres are today strictly entangled and individual and collective human behavior is increasingly predictable from, or even influenced by, the dynamics of the digital world.

The pivotal role that social media had in coordinating the protesters in the 2011 Arab Spring wave (Panisson 2011), the effectiveness of online communication on political opinion formation (Conover et al. 2010), or the rapidity of spreading of earthquakes alarms through social media (Sakaki, Okazaki, and Matsuo 2010) are clear examples of the direct effect that the Web can have on human lives.

Mining the user-generated data streams from social media allows the analysts to detect, better understand and even try to predict real world events such as the spreading of diseases (Culotta 2010), the mood of the inhabitants of urban

areas (Quercia et al. 2010), the trends in the stock market (Bollen, Mao, and Zeng 2011), or even the political election results (Tumasjan et al. 2010; T. Metaxas, Mustafaraj, and Gayo-Avello 2011). The social and economic relevance of such opportunities has increased the interest not only in learning online social dynamics more in depth but also in understanding the extent to which such dynamics can be manipulated and controlled. For instance, Twitter has been leveraged to unfairly boost the consensus on some politicians for electoral purposes by automatically generating high amounts of fake tweets supporting particular opinions (Ratkiewicz et al. 2011).

Efforts have been spent in developing bot-detection techniques (Lee, Caverlee, and Webb 2010), but the real potential of robots in manipulating online social networks is still widely unknown. We contribute to the exploration of this field by investigating the relation between *trust*, *popularity* and *influence*, three building blocks on which dynamics of opinion formation and information spreading on social substrates are based. Differently from previous data-driven studies on this matter, we go beyond the exploratory data analysis and we set up a social experiment based on a *bot* interacting with human users in an online social environment. Our goal is to understand to what extent a bot with no public information can become popular and influential or, more in general, how it can impact the dynamics of the network.

Surprisingly, our experiments reveal that an untrustworthy individual can become very relevant and influential through very simple automated activity. Moreover, the social reactions aroused by the bot's activity give interesting insights on the possibility of automated attacks to affect the structure of the social substrate.

Contributions and roadmap

Obtained results 1) provide explanations on the roots of *popularity* and *influence* in a social media context, 2) give useful insights about how a social *recommender* (or *spammer*) can be extremely effective, 3) provide a direct test of accuracy of modern *recommendation* techniques on a real user base, and 4) show how the intervention of an external, anomalous entity in a social network can alter its dynamics and unveil *social polarization* phenomena. Moreover, differently from traditional data-driven analysis that must face the hard task of inferring causal relationship between facts just by ob-

	Jul 2010		Nov 2011	
	Social	Comm	Social	Comm
Nodes	86,800	72,054	179,653	95,121
Links	697,910	483,151	1,566,369	677,652
$\langle k_{out} \rangle$	8.0	6.71	8.72	7.12
Reciprocation	0.57	0.59	0.51	0.60
GSCC size	62,195	33,700	127,997	44,072

Table 1: Basic structural quantities in the aNobii social and communication networks at the beginning and at the end of the experiment. $\langle k_{out} \rangle$ is the average out-degree, reciprocation is the ratio of bidirectional arcs, and GSCC is the greatest strongly connected component.

serving their temporal sequence, we can directly explore the causality of events by acquiring the perspective of an agent that interacts with other users. We make available upon request the anonymized data collected during and right after the experiment.

In Section 2 we describe the main features of the social media on which we based our experiment and we present the structure of the bot we used to interact with other users. In Sections 3 and 4 we explain the results obtained by the bot in terms of gain of popularity and influence, and in Section 5 we overview the side effects of its activity in terms of altering the dynamics of interaction on the social network.

2 Dataset and experimental setup

Our experiment has been set up on aNobii.com, a social network for book lovers. Besides specifying classic profile information such as age and geographic location, aNobii users can fill a personal digital *library* with the titles they have read and they can enrich them with ratings, annotations and reviews. Users socialize by means of thematic groups and by establishing pairwise social connections. Social ties are directed and can be created without any consent of the linked user. Two types of mutually-exclusive social links are available, namely *friendship* and *neighborhood*, but their behavior is identical, therefore in the following we consider a single *social* network composed by their union. Every profile page contains a public *shoutbox*, where any user can write messages. Message exchange implies the existence of another social graph that we call *communication* network, where a directed arc $A \rightarrow B$ models a message from A to B 's shoutbox. The arcs are weighted with the cumulative number of messages. Even though a few privacy settings are available, the whole profile information is usually public.

Table 1 reports basic structural statistics of social and communication graphs at the beginning and at the end of the 1.5 years time frame of our experiment. The network roughly doubled its size in that period and kept a high reciprocation trend in social linking and a rather low average degree if compared to other general-purpose social media (Ugander et al. 2011). However, not surprisingly, the social activity in terms of social linking and message exchange is rather broad, as shown by Figure 1.

We started collecting data in September 2009 through a periodic activity of Web crawling to study the dynamics of




Figure 1: Complementary cumulative density functions of in- and out-degree $k_{in/out}$ in the social and communication graphs and of in- and out-weighted degree $msg_{in/out}$ (i.e., overall number of received/sent messages) in the communication graph at the end of year 2011.




Figure 2: Average number of books and sent messages for users who received the same number of messages msg_{in} .

link creation in the network (Aiello et al. 2010). The relatively small size of the graph allowed us to crawl all the nodes reachable by a BFS strategy once every 15 days. We were constrained to create an empty stub account to access some protected sections of user profiles accessible only to subscribed users, and use its credentials to access such pages. The profile was automatically named *lajello* after the email prefix of the subscriber.

Our experiment started with an accidental discovery. In July 2010 the default user settings of the website silently changed so that every visit of a logged user automatically started leaving a trace in a private *guestbook* of the visited profile. The users can check their personal guestbook and see the last 30 visitors of their libraries. As a result, our crawler left a trace of its passage in all the profiles reached by the BFS approximatively twice a month.

The unexpected reactions the bot caused by its visits motivated us to set up a social experiments in two parts to answer the question “*can an individual with no trust gain popularity and influence?*”. In the first part we kept running the periodic crawling to check the variation over time of the user response to the bot visits. In the latter phase we run a modified version of our bot that was able to communicate with other users, asking them to take some action. Results are discussed in Sections 3 and 4.

3 Gain of popularity

The popularity of an individual in social media is widely considered to be strictly related to the amount of incom-




Figure 3: User response to *lajello*'s probing activity taking place between Jul 2010 and Jan 2012. The timeline at the top represents the overall histogram of the number of received messages on a daily basis during the whole period. The times of start and stop of the periodic crawling activity are marked. Bottom-left plot shows a finer-grain histogram for the period from Jul 2010 to Mar 2011. The cumulative number of distinct users writing messages is reported as well. Bottom-right plot shows a zoomed view on the 36 hours after the first recommendation message sent by the bot (see Section 4).

ing connections or to the number of received messages, meaning that a popular user is characterized by a network structural position at the center of star-shaped clusters (Stoica, Couronné, and Beuscart 2010). Researchers have long investigated the defining traits of the popular users in social media. Factors determining user popularity range from the amount and the kind of profile features exposed to the public (Lampe, Ellison, and Steinfield 2007; Strufe 2010) to the extroverted attitude of individuals in social communication (Quercia et al. 2012). More in general, the creation of social trust bonds is strongly correlated with the similarity of opinions and profiles (Golbeck 2009).

From a quantitative point of view, also in aNobii we find a clear positive correlation between the number of incoming messages with the amount of books publicly shown in the library and with the amount of sent messages, as shown in Figure 2. This is not surprising given the tendency to reciprocity of social protocols and considering that aNobii is focused on books only, but still it gives an hint about the importance of being *active* to gain the attention of other participants. However, identifying the main cause of popularity among many user activity indicators that are all strictly correlated one to another is not trivial. In this perspective, our social experiment is aimed to isolate a single, minimal social activity (namely the visit to other profile pages) and to test its effectiveness in gaining popularity.

As soon as the new aNobii settings on visits took place in July 2010, our bot started collecting a massive response from visited users, and every round of visits triggered a burst of comments on the bot's public wall. Figure 3 shows the histogram of the number of visits in time. Higher spikes occurs right after the start of the two-monthly visits. The different height of spikes depends from the duration of the crawl, that has been inconsistent (from 24 hours to some days) due to the discontinuous uptime of the Web service. We also ob-




Figure 4: Complementary cumulative density functions of Pagerank and HITS in social and communication networks. A guide to the eye marks the upper bound of the 99th percentile. The scores of *lajello* are marked with blue crosses and belong, from left to right in the figure, to the 98th, 99th, 98th, and 96th percentiles.

serve that the cumulative number of distinct users shouting grows constantly and roughly linearly in time.

Interestingly, the rapid bursts of messages happening after the visits do not trigger any self-feeding effect in the incoming communication stream. When the crawler terminates the visit round, people quickly stop sending messages except to respond with the same feedback when next round is performed. To confirm this, in mid February, after 15 crawling rounds, we stopped the bot until September 2011. A negligible amount of messages has been received in the inactivity period and the user feedback recovered the same pace when we resumed the crawling.

At the end of its activity, *lajello*'s profile has become one of the most popular in the website, boasting 2,435 public messages from 1,263 different users, more than 200 private messages, more than 66,000 visits to its profile, and 211 incoming social connections (partitioned in 125 neighbors and

Feature	Description
Common neighbors	$CN(u, v) = \Gamma_{out}(u) \cap \Gamma_{in}(v) $
Triangle overlap	$\frac{CN(u, v)}{\Gamma_{out}(u)}$
Resource allocation	$\frac{1}{k_{out}(u)} \sum_{z \in (\Gamma_{out}(u) \cap \Gamma_{in}(v))} \left(\frac{1}{k_{out}(z)} \right)$
Reciprocation	Binary attribute, whether the inverse link (v, u) is already present
Contact list	σ_{sim} between social contacts vectors
Library	σ_{sim} between library vectors
Groups	σ_{sim} between group vectors
Group size	Size of the smallest group the two users have in common

Table 2: Features used in the recommendation of a directed link between users u and v , along with their description. $\Gamma_{in/out}(u)$ is the set of u 's in/out neighbors, $k_{out}(u) = |\Gamma_{out}(u)|$. Resource allocation has been defined by Zhou et al. 2009. σ_{sim} denotes the cosine similarity.

86 friends). Considering the very limited size of the social network and the range of the degree distributions of aNobii users (see Figure 1), such values are extremely high. In particular, in December 2011, the bot was the 2nd user for the number of distinct people writing on its shoutbox (the first one is a well-known writer) and the 28th in the number of received messages, with just 10 months of activity. Moreover, it was among the 0.3% of users with the highest in-degree.

Even though the amount and weight of incoming connections is usually considered a good proxy for popularity, more sophisticated metrics taking into account the social status of the connected nodes could give a more accurate picture on the amount of popularity gained, following the intuition that an individual is more popular if it is followed by other popular people. To this end, we computed the distributions of Pagerank and HITS on social and communication networks (Figure 4). Even considering such metrics, *lajello* positions itself in the very tail of the distributions.

In short, our experiment gives strong support to the thesis that popularity can be gained just with continuous “social probing”. This is particularly striking from the perspective of the security of online social environments since we have shown that a very simple *spambot* can attract great interest even without emulating any aspects of a typical human behavior.

4 Influence

Previous studies on the interdependence between popularity and influence in social media show that popularity does not imply the ability of being influential and vice-versa (Cha et al. 2010). It appears that influence depends on a combination of the position on the social graph (Ilyas and Radha 2011) and of the language used in public communications (Quercia et al. 2010) and, intuitively, trust seems to be a precondition for both popularity and influence.

However, the great response collected from the bot with a very simple automated activity naturally arises the question about the possibility for an untrusted, empty profile to be influential on other community members.

To answer this question we arrange a second phase of the




Figure 5: Logical scheme of our recommender system. The target user u together with its distance-2 neighbors and the respective features are given to the classifier. The best match is then given in output.

social experiment in which the bot tries to convince users to perform an action. Many possible activities are allowed by the aNobii service, but we focus on the task of persuading users to add a new neighbor in their contact list. This choice is motivated by two reasons mainly. First, recommending a social activity rather than one that you can carry on individually (such as the adoption of a new book) has a higher potential of triggering some interaction between the users that are subjected to the experiment. This is desired since we are also interested in studying the social and emotional reactions of the user base to the bot’s activity. Second, we observe that despite the vast literature on link recommendation in social networks (Lü and Zhou 2011), as far as we know there are no reported results on the actual effectiveness of the most recent suggestion techniques on actual users. Our experiment is structured such that the validity of modern contact recommendation techniques and the bot’s degree of influence can simultaneously be validated.

First, let us introduce our link recommendation method. Based on the findings of previous work on link prediction and recommendation (Aiello et al. 2012) we adopt a machine-learning approach that combines features extracted from the profile pages and structural features defining the position of users on the social graph. Predictive metrics usually exploit the observation that the process of link creation in social networks is driven by proximity and homophily (Aiello et al. 2010) i.e., people tend to close social triangles with people more similar to them.

We use the Rotation Forest classifier—which proved to be the best performing for this task among all WEKA classifiers (Hall et al. 2009)—trained on the feature set listed in Table 2. Among a larger set of features, we found those to be the most predictive through a Chi Squared feature selection. To first test the effectiveness of the classifier we used a balanced training set of 20k pairs of users residing 2-hops away in the social graph. Positive examples are pairs that got connected with a social tie between September and November 2010 and negatives are chosen randomly among those who did not. On a 10-fold cross validation we obtained 0.82 as Area Under the ROC Curve (AUC), which denotes a good quality of the selected predictors.

The trained classifier can be easily used to produce recommendations on the fly. Given a user u , the classifier is ap-

plied between u and all her distance-2 neighbors Γ_{out}^2 . All the contacts in Γ_{out}^2 for which the classifier gives a positive response are sorted by the classifier’s confidence score and the contact at the top of the list is selected as the one to recommend. The logical scheme of the recommender is depicted in Figure 5. We consider only distance-2 neighbors to keep the computational complexity low, but this strategy does not worsen much the quality of the recommendation given that in the most frequent case links are created between distance-2 pairs (Aiello et al. 2010).

Once our recommender has been set, we used it to perform direct recommendations to the aNobii users. Differently from the first phase of the experiment, we focused on users who declared Italian nationality on their profile. This is motivated by the fact that 60% of the aNobii users are from Italy and they represent the most active core of the network.

We structured the experiment as follows. First, we listed all the Italian users that wrote at least one message on the bot’s shoutbox (we call them *followers*) and an equal number of those who did not (*non-followers*) taken at random among all the other Italian users with at least 10 books in their library, to avoid inactive users. Then we randomly sampled 50% of users from this pool and produced for them a single contact suggestion using our recommender. For the remaining 50% we produced random recommendations. Finally, we randomly sampled 25% of the recommendation pairs (u, v) (where u is the user to be recommended and v the recommended contact) and we produced the corresponding reciprocal recommendations (v, u) . At the end of this procedure we obtained five equal-size recommendation categories (followers recommendation, followers random, non-followers recommendation, non-followers random, and reciprocal) for a total of approximately 3,000 different recommendations of non-existent social ties.

Using a random message generator that has been written specifically for this experiment, we produced a message in Italian language for each user in the recommendation sample. Messages were all different for number and type of sentences used and for the use of lexical classes but all of them contained a suggestion for the recipient to add a certain user to her neighbors list.

We sent all the recommendation messages to the public shoutboxes of target users during the night of November 8th 2011. From the early morning after we registered a bursty user response. In the 36 hours following the experiment, more than 400 shouts were published on *lajello*’s shoutbox. Histogram in Figure 3 shows how the peak of messages received during the day after is roughly one order of magnitude higher than any other peak observed during the past crawling activity.

An even more surprising feedback was given by the link recommendation results. Among the 361 users who created at least one social connection in the 36 hours after the recommendation, the 52% followed the suggestion given by the bot. The distribution of the successful suggestions over the different categories (Figure 6) reveals three interesting patterns.

First, non-random recommendations dispensed to followers have been far more effective than those given to non-




Figure 6: Distribution of successful contact suggestions among the different recommendation classes. “Followers” denote the users who wrote at least one message to *lajello*. “Recommend” and “random” indicate whether the recommendation has been performed using our classifier or just picking a random user. “Reciprocal” class includes recommendations of links that reciprocates other recommendations taken at random from any of the previous classes.

followers, in other words, *lajello* has a greater persuasive power over those who are more aware of its presence and activity. Second, the effectiveness of thoughtful suggestions based on our recommender is at least double than random recommendations, meaning that users are mainly reluctant to obey blindly (even though a non-negligible 11% of the successful recommendation belongs to the random categories). Finally, it is surprising to see that recommendations are more likely to succeed when both directions of the link are suggested to the two endpoints. This is mainly due to the fact that bidirectional recommendations are more likely to trigger some communication between the two endpoints that in many cases ends up in the agreement to form a new social tie.

Shortly after registering these results, the bot’s account was suspended by the system administrators and all the recommendation messages were canceled. Even though we could not register the feedback of all the recommended users, the obtained results still lend themselves to meaningful interpretations from two different points of view.

From a social network analysis perspective, we observe that untrustworthy individuals in social media can become influent if they have enough popularity, which can be gained with automated activity. Even if our experiment has been forced to be limited in time and restricted to a relatively small number of users, its outcomes give a very strong evidence about the extent to which a community could be manipulated by external agents. We consider this a surprisingly outstanding result in this field.

From the service provider angle instead, we learn that a *social recommender* can be very effective (52% of success ratio in our case). Even if fair comparisons across different recommendation techniques and datasets are difficult, our results appear to be even better than recent results on link recommendation based on temporal datasets (Backstrom and Leskovec 2011). This suggests that many recommendations that are labeled as false positives in such data-driven studies may result to be successful when they are ac-




Figure 7: Stacked area graph depicting the ratio of messages with positive (bottom blue), negative (top red), or neutral (center green) connotation in *lajello*’s shoutbox. On every point on the x-axis we present the statistic over 50 consecutive messages.

tually proposed to the end user. Unfortunately, the same argument can be applied to malicious attackers: external unauthorized robots programmed to pilot user opinions and actions can be extremely effective in a social media context. For this reason, additional care should be put in the design of security devices of social media to avoid such hazards.

5 Emotional reactions and social polarization

In previous Sections we focused on the quantitative analysis of the user response to our bot to assess its popularity and influence across the community, but an additional relevant outcome of the experiment pertains the qualitative sphere of the collected user reactions.

The most striking result that is perceivable at a first glance is the wide variety of sentiments expressed in the messages received by *lajello*, ranging from veneration to hatred. This has been allowed by the fact that many people believed the bot to be a human user and only in the latest phase of the experiment rumors about the possibility of *lajello* being a robot began to slowly spread over the network. This result is coherent with another recent “socialbot” experiment on Facebook (Boshmaf et al. 2011), in which 80% of unaware users accepted the friendship request of a bot, as well as with previous experiments on social phishing (Jagatic et al. 2007).

Moreover, many users revealed deep discomfort for the fact that an unknown user was periodically visiting their profiles, even if our bot explored just the information available to anyone else in the community. This empirically confirms previous studies on the perception of privacy of users on online social media (Levin and Abril 2009) showing that users are very concerned about the privacy risks given by the requests of *strangers* without a good reputation or with no profile picture. Moreover, another explanation for such emotional reactions is given by the *fear of control*, widely discussed by Michel Foucault (1975), who exemplifies this phenomenon using the paradigm of the *Panopticon*, the architectural model of a structure in which people actions are constantly observed and recorded. Several users openly expressed in their messages the fear of control or the perception of a strong privacy violation.

A fine-grain text analysis that is able to capture all the sentiment nuances expressed by the users goes beyond the




Figure 8: Networks induced by the members of the two groups opposing on *lajello* topic. Blue and red circles represent users supporting and criticizing the bot, respectively. Node size is proportional to the degree. Singleton nodes, directionality and weight of edges are not shown for simplicity.

	Actual		Randomized	
	Intra	Inter	Intra	Inter
Social	74%	26%	52%	48%
Communication	71%	29%	49%	51%

Table 3: Portion of links that reside inside a cluster (intra) or, conversely, that connect users belonging to different clusters (inter) in the real networks and in their randomized versions.

scope of this work. However, to assess the amount of users that were positively or negatively impressed by the activity of the bot, we manually annotated all the received messages with a ternary label *positive/neutral/negative*. All the annotations have been performed by three authors and a majority vote has been applied to pick the final label. Just a few ambiguous cases occurred. On average, we found that 59% of messages have a positive connotation and 16% are negative instead. By expanding this aggregate data on a temporal scale we note interesting fluctuations in the trend of consensus and dissensus, as shown in Figure 7.

Initially, users notice for the first time the visits of the bot and reply giving it a warm welcome. Soon after, the unknown reason of the visits and the permanently empty profile start raising some uncertainty in the community, and people begin asking for explanations but without expressing an explicit sentiment. Meanwhile, a small but growing percentage of hostile messages were addressed to the bot. The number of neutral messages gets thinner in time, even if the positive mood remains always prevalent. The most evident polarization of opinions occurs in the last 400 messages, after the recommendation, when people started discussing whether the behavior of *lajello* could have been acceptable. This phenomenon is coherent with studies on social polarization on Twitter (Yardi and boyd 2010), where user interactions on controversial topics reinforce the separation of opposing opinions.

Polarization in user opinions could be also tracked from

the analysis of thematic groups. We counted seven groups that have been spontaneously dedicated to *lajello* since its appearance. The main threads of discussion were focused on exchanging opinions about its identity, but also strong judgments about its activity have been expressed. In particular, right after receiving the recommendations from the bot, two opposing groups were quickly created: one in favor of *lajello*, named “Nessuno tocchi lajello” (tr. “Hands off lajello”) currently with 102 members, and one against it with 72 participants, named “I consigli (non richiesti) di lajello” (tr. “The (non-requested) suggestions of lajello”).

The sharp difference of opinions in the two groups motivated us to investigate how much the two factions are clustered when considering their social connections. Figure 8 depicts the social and communication networks induced by the members of the two groups, displayed according to a force-directed layout (Jacomy et al. 2011). A first visual inspection already reveals that the two node categories tend to form two distinguishable clusters, even in a state of high link density.

To verify the presence of separate clusters more precisely, we perform a randomization test. We count the ratio between edges that fall within the same cluster (also called *intra-cluster edges*) and those that connects nodes in different clusters (*inter-cluster edges*) and we compare the obtained values with the same measurement performed on a randomized version of the network where nodes keep their out-degree but their links are rewired at random. Results in Table 3 reveal a much stronger tendency of nodes to be connected inside their own cluster than in a random model.

An additional evidence of user polarization into two opinion clusters can be given by applying community detection algorithms to the networks and check whether the detected clusters match the group membership of nodes (Conover et al. 2011). We perform this test using the OSLOM community detection algorithm (Lancichinetti et al. 2011) due to its good performance over other state of the art network clustering techniques. Since OSLOM outputs more than two clusters and we are interested in a binary polarization, we greedily merge together clusters whose majority of nodes belong to the same group, until we obtain two macro-clusters.

To assess the adherence of clusters to groups we compute the average fraction of correctly classified vertices (Girvan and Newman 2002) over 1000 realizations of OSLOM with random seed. We find that 80% and 72% of nodes in the social and communication networks respectively are correctly classified. This confirms the intuition that the emergence of opposed factions was rooted in the structure of social connections, but has been hidden until the system perturbation caused by *lajello*. As shown in classical studies on social networks (Zachary 1977), we can interpret this result also as the possibility of disrupting the network along the edges of detectable clusters, in a *divide-et-impera* logic, by introducing a perturbing element in the community.

Finally, besides motivating the creation of new groups, the bot activity triggered also the creation of new communication links and new nodes in the network. Users started to talk about *lajello* not only through the mentioned groups but also posting messages on each other shoutboxes. Tracking the keyword “*lajello*” we can extract a subset of the com-




Figure 9: Small subsample of communication network where edges represent messages containing the keyword “*lajello*”. Nodes are resized and colored according to their degree. The big red node on the right is a user who is emulating the behavior of *lajello* to become popular. In this snapshot it is starting to attract the interest of other users who are writing on its shoutbox.

munication network that maps the interaction channels over which our bot has been mentioned. The analysis of such filtered network reveals the presence of some nodes with an in-degree sensibly higher than the average, like the one shown in Figure 9. Such nodes are *clones* of our bot, i.e., empty profiles created to silently visit other users in the same way *lajello* does. The activity of the clones soon attracted the attention of other users, who started writing messages on their walls mentioning *lajello*.

6 Conclusions

We presented the outcome of a social experiment aimed to explore the relation between *trust*, *popularity* and *influence* in the dynamics of online social media. We showed that popularity in social networks does not require peculiar user features or actions, since an automated agent can acquire a surprisingly high popularity just by reiterating a simple activity of “social probing”. In a second phase of the experiment we sent friendship suggestions from our bot to a pool of users, providing random suggestions to some and thoughtful recommendations to others. As evidence that an untrustworthy user can be very influential if popular enough, we found that people more aware of the presence of the bot have been more inclined to follow its suggestions. Moreover, unlike previous work in the field, we validate the effectiveness of modern link recommendation techniques with an explicit feedback of the user base. Last, we registered profound social dynamics alterations. Specifically, our bot was widely mistaken as a human user and its activity triggered the creation of new groups and communication channels, originated emulous users, raised explicit concerns about privacy violations and aroused a diffused fear of being controlled. Moreover, it unveiled preexistent latent social network clusters by triggering the polarization of opinions of community members.

Our results contribute to shed light on the effects that of automated agents can have in online social environments. Our findings have important implications on the way in which we conceive security and privacy in social networks. Since sociological and emotional sides of the experiment are crucial to investigate this issues, we plan for the future a

more systematic sentiment analysis of the messages originated by the users in the wake of the bot's activity.

Acknowledgments

This work is funded by the projects "INCA" (Finpiemonte) and "Information Dynamics in Complex Data Structures" (PRIN). We are grateful to aNobii for making its data available, to Giuseppe Tipaldo, Alain Barrat, Ciro Cattuto and to the members of the NaN group at Indiana University for helpful discussions. Graphs visualizations are made with Gephi (Bastian, Heymann, and Jacomy 2009).

References

- Aiello, L. M.; Barrat, A.; Cattuto, C.; Ruffo, G.; and Schifanella, R. 2010. Link creation and profile alignment in the aNobii social network. In *SocialCom'10*, 249–256.
- Aiello, L. M.; Barrat, A.; Schifanella, R.; Cattuto, C.; Markines, B.; and Menczer, F. 2012. Friendship prediction and homophily in social media. *ACM TWEB*. To appear.
- Backstrom, L., and Leskovec, J. 2011. Supervised random walks: predicting and recommending links in social networks. In *WSDM'11*, 635–644.
- Bastian, M.; Heymann, S.; and Jacomy, M. 2009. Gephi: an open source software for exploring and manipulating networks. In *ICWSM'09*. AAAI.
- Bollen, J.; Mao, H.; and Zeng, X. 2011. Twitter mood predicts the stock market. *Journal of Comp. Sci.* 2(1):1–8.
- Boshmaf, Y.; Muslukhov, I.; Beznosov, K.; and Ripeanu, M. 2011. The Socialbot Network: When Bots Socialize for Fame and Money. In *ACSAC'11*.
- Cha, M.; Haddadi, H.; Benevenuto, F.; and Gummadi, K. P. 2010. Measuring User Influence in Twitter: The Million Follower Fallacy. In *ICSWM'10*.
- Conover, M. D.; Gonçalves, B.; Ratkiewicz, J.; Flammini, A.; and Menczer, F. 2010. Predicting the Political Alignment of Twitter Users. In *SocialCom'11*.
- Conover, M.; Ratkiewicz, J.; Francisco, M.; Gonçalves, B.; Flammini, A.; and Menczer, F. 2011. Political Polarization on Twitter. In *ICSWM'11*.
- Culotta, A. 2010. Towards detecting influenza epidemics by analyzing Twitter messages. In *SOMA'10*, 115–122. New York, NY, USA: ACM.
- Foucault, M. 1975. *Discipline & Punish: The Birth of the Prison*. Random House.
- Girvan, M., and Newman, M. E. J. 2002. Community structure in social and biological networks. *PNAS* 99(12):7821–7826.
- Golbeck, J. 2009. Trust and nuanced profile similarity in online social networks. *ACM TWEB* 3:12:1–12:33.
- Hall, M.; Frank, E.; Holmes, G.; Pfahringer, B.; Reutemann, P.; and Witten, I. H. 2009. The WEKA data mining software: an update. *SIGKDD Explorations* 11:10–18.
- Ilyas, M. U., and Radha, H. 2011. Identifying Influential Nodes in Online Social Networks Using Principal Component Centrality. In *ICC'11*, 1–5.
- Jacomy, M.; Heymann, S.; Venturini, T.; and Bastian, M. 2011. ForceAtlas2, A Graph Layout Algorithm for Handy Network Visualization. Technical report, Gephi Consortium.
- Jagatic, T. N.; Johnson, N. A.; Jakobsson, M.; and Menczer, F. 2007. Social phishing. *Comm. of ACM* 50:94–100.
- Lampe, C. A.; Ellison, N.; and Steinfield, C. 2007. A familiar face(book): profile elements as signals in an online social network. In *CHI'07*, 435–444. New York, NY, USA: ACM.
- Lancichinetti, A.; Radicchi, F.; Ramasco, J. J.; and Fortunato, S. 2011. Finding Statistically Significant Communities in Networks. *PLoS ONE* 6(4).
- Lee, K.; Caverlee, J.; and Webb, S. 2010. Uncovering social spammers: social honeypots + machine learning. In *SIGIR'10*, 435–442. New York, NY, USA: ACM.
- Levin, A., and Abril, P. S. 2009. Two Notions of Privacy Online. *Vanderbilt JETLaw* 11:1001–1051.
- Lü, L., and Zhou, T. 2011. Link prediction in complex networks: A survey. *Physica A* 390(6):1150–1170.
- Panisson, A. 2011. Visualization of Egyptian Revolution On Twitter. <http://slashdot.org/story/11/02/15/1544254>.
- Quercia, D.; Ellis, J.; Capra, L.; and Crowcroft, J. 2010. In the Mood for Being Influential on Twitter. In *SocialCom'11*.
- Quercia, D.; Lambiotte, R.; Kosinski, M.; Stillwell, D.; and Crowcroft, J. 2012. The Personality of Popular Facebook Users. In *CSCW'12*.
- Ratkiewicz, J.; Conover, M.; Meiss, M.; Gonçalves, B.; Flammini, A.; and Menczer, F. 2011. Detecting and Tracking Political Abuse in Social Media. In *ICSWM'11*.
- Sakaki, T.; Okazaki, M.; and Matsuo, Y. 2010. Earthquake shakes Twitter users: real-time event detection by social sensors. In *WWW'10*, 851–860. New York, NY, USA: ACM.
- Stoica, A. M.; Couronné, T.; and Beuscart, J.-S. 2010. To Be a Star Is Not Only Metaphoric: From Popularity to Social Linkage. In *ICSWM'10*.
- Strufe, T. 2010. Profile popularity in a business-oriented online social network. In *SNS'10*, 2:1–2:6.
- T. Metaxas, P.; Mustafaraj, E.; and Gayo-Avello, D. 2011. How (Not) to Predict Elections. In *SocialCom'11*, 165–171.
- Tumasjan, A.; Sprenger, T.; Sandner, P.; and Welp, I. 2010. Predicting elections with Twitter: What 140 characters reveal about political sentiment. In *ICSWM'10*, 178–185.
- Ugander, J.; Karrer, B.; Backstrom, L.; and Marlow, C. 2011. The Anatomy of the Facebook Social Graph. *CoRR* abs/1111.4503.
- Yardi, S., and boyd, d. 2010. Dynamic debates: An analysis of group polarization over time on Twitter. *Bulletin of Science, Technology and Society* 20:S1–S8.
- Zachary, W. W. 1977. An information flow model for conflict and fission in small groups. *Journal of Anthropological Research* 33:452–473.
- Zhou, T.; Lü, L.; and Zhang, Y.-C. 2009. Predicting missing links via local information. *European Physical Journal B* 71(4):623–630.