

BFF: A Tool for Eliciting Tie Strength and User Communities in Social Networking Services

Ricard L. Fogués¹, Jose M. Such², Agustin Espinosa¹ and Ana Garcia-Fornes¹

Received: date / Accepted: date

Abstract The use of social networking services (SNSs) such as Facebook has explosively grown in the last few years. Users see these SNSs as useful tools to find friends and interact with them. Moreover, SNSs allow their users to share photos, videos, and express their thoughts and feelings. However, users are usually concerned about their privacy when using SNSs. This is because the public image of a subject can be affected by photos or comments posted on a social network. In this way, recent studies demonstrate that users are demanding better mechanisms to protect their privacy. An appropriate approximation to solve this could be a privacy assistant software agent that automatically suggests a privacy policy for any item to be shared on a SNS. The first step for developing such an agent is to be able to elicit meaningful information that can lead to accurate privacy policy predictions. In particular, the information needed is user communities and the strength of users' relationships, which, as suggested by recent empirical evidence, are the most important factors that drive disclosure in SNSs. Given the number of friends that users can have and the number of communities they may be involved on, it is infeasible that users are able to provide this information without the whole eliciting process becoming confusing and time consuming. In this work, we present a tool called Best Friend Forever (BFF) that automatically classifies the friends of a user in communities and assigns a value to the strength of the relationship ties to each one. We also present an experimental evaluation involving 38 subjects that showed that BFF can significantly alleviate the burden of eliciting communities and relationship strength.

Keywords: Information retrieval, social network, social media, privacy, tie strength.

¹Departamento de Sistemas Informáticos y Computación, Universitat Politècnica de València, Spain

{rilopez,aespinos,agarcia}@dsic.upv.es

²School of Computing and Communications, Lancaster University, UK

j.such@lancaster.ac.uk

1 Introduction

Social networking services (SNSs) are currently the services that are most more demanded by users worldwide. Facebook (with more than 800 million active users¹) and Flickr (with 51 million registered members²) are two of the most successful SNSs. People register to these SNSs and share images, videos, and thoughts because they perceive a great payoff in terms of friendship, jobs, and other opportunities [5]. However, the huge number of items uploaded to these SNSs and the persistence of these items in the social networks have the potential to threaten the privacy of their users[11]. For example, employers are becoming accustomed to checking the profile of the candidates in popular SNSs. If the privacy of the profile of a candidate is not properly set, what an employer sees in that candidate's profile may affect the employer's decision. It might even be possible for a stalker to infer the address of a person by looking at that person's photos posted in a social network.

Factors like the increase in public attention to privacy matters and the users' increment of familiarity with SNS technology have increased the privacy concerns in SNSs [2]. To cope with privacy threats users tend to adjust and modify the default privacy preferences set up by the SNSs since they feel that these default settings are not enough to protect them. Nonetheless, the current privacy setting mechanisms offered by SNSs seem difficult or confusing for users [17][27]. These complications and obstacles lead to privacy policies that do not fit users' preferences, and, in turn, discourage users to show high engagement in terms of how much they participate in the SNS (e.g., the amount of photos they upload) [26].

To address this, privacy management mechanisms that are able to automate the process of privacy policy definition as much as possible are needed [6]. In this way, our long-term aim is to develop intelligent software agents that could act as privacy assistants recommending adequate privacy policies. To this aim, we need to consider the existing empirical evidence on what drives disclosure in SNSs. In particular, it has been proven that the most important factors that users consider to decide whether disclosing information is adequate or not are the strength of the relationships they have to others and the communities in which they are involved [29]. Thus, an intelligent privacy assistant agent should base its predictions on this information. The problem is that eliciting this information from the user could become a time consuming process, e.g., it would require that users specify for each of his/her friends how strong their relationship is (average user in Facebook has 130 friends [21]).

As introduced in the paper of Granovetter[10], the concept of *tie strength* defines the relationship between two individuals. In his work, Granovetter describes two different types of ties: *strong* and *weak*. On the one hand, strong ties usually include relationships such as family and close friends. On the other hand, weak ties may refer to coworkers or less trusted friends. More

¹ Facebook statistics <http://newsroom.fb.com/>

² Yahoo advertising solutions <http://advertising.yahoo.com/article/flickr.html>

recent works have proposed models to predict tie strength in SNSs [8, 13, 30]. These works showed that it is possible to infer tie strength from the available personal data in a SNS. However, these works were not aimed at creating an actual tool. Thus, they only considered the predictive capabilities of the variables collected from the SNS without taking into account other factors that are of crucial importance to create an usable tool that predicts tie strength. In particular, these works did not consider either: (i) the computational cost of collecting the variables from the SNS (e.g., if the tool takes too much to complete the process this could be also seen as a time-consuming and not feasible in practice); or (ii) if all the variables apply to any possible Facebook user (e.g., these works consider language-dependent variables that could limit the applicability of their approach to speakers of other languages).

Regarding communities, they are usually defined as natural divisions of network nodes into densely connected subgroups [9]. In our context, the nodes are the contacts or friends of a given participant, and the connections between the nodes are friend relationships. Although many SNSs support some notion of community by means of allowing groups of users, they usually require that the user manually assigns each friend to the corresponding group. For example, Facebook offers the possibility of creating groups of friends, and then assigning privacy policies for each of them. However, this again has the problem that users are required to spend a considerable amount of time creating the groups and assigning friends for each of the groups (if we consider that the average number of friends in Facebook is 130, classifying all of them into groups can represent a serious challenge). Thus, the process of friend grouping should be also automatized as much as possible. To this aim, we can use one of the many existing community finding algorithms [7]. The problem is that, to the best of our knowledge, there is no empirical evidence of how this community finding algorithms perform when they are applied to real social graphs extracted from a SNS, so choosing the most appropriate one is a challenging problem.

Our main contributions in this article are:

1. We present a new tool called Best Friend Forever (BFF) that is able to automatically obtain relationship strength values and user communities from a SNS. Moreover, it allows users to further refine the results if they are not accurate enough. BFF has been implemented as a Facebook application and is publicly available at gti-ia.dsic.upv.es/bff.
2. We propose a new method to calculate tie strength in SNSs that considers not only the predictive capability of the variables used but also other crucial factors to develop an actual tool: the temporal cost of collecting the variables and that these variables are general enough to be applied for any SNS user. Moreover, this new method has been implemented in the tie strength module of BFF.
3. We evaluate several community finding algorithms using real social graphs (from 38 real Facebook users), and select the most appropriate one to be included in the community finding module of BFF based on their accuracy and their temporal cost.

4. We empirically demonstrate that by using BFF we are able to elicit users' relationship strength and communities requiring little intervention from users. In particular, 81.71% of tie strength values were exactly inferred and 67.08% of users' friends were correctly organized into communities.

The rest of the paper is organized as follows. Section 2 presents an overview of BFF and its different elements. Section 3 explains the tie strength prediction module and how it works. Section 4 presents the community prediction module and the community finding algorithms used by this module. Section 5 reports the results of the experimental evaluation and discusses its generalizability and limitations. Section 6 discusses some related works. Finally, Section 7 concludes the paper and outlines future research directions.

2 Best Friend Forever

This section introduces our tool and gives a complete overview of it. BFF aims to retrieve information from the social network that can be useful to recommend privacy policies. Specifically, the data needed is tie strength and friend groups. BFF is written in PHP and Javascript and is publicly accessible. Due to our experimentation needs, BFF is currently working as a web page; however, in the future, we plan to distribute BFF as a software program that users can execute in their own computers or on a trusted web server in order to preserve their privacy.

BFF is composed of two modules: (i) community prediction, and (ii) tie strength prediction. The community prediction module is in charge of creating chunks of users from the participant's contacts. The tie strength prediction module establishes a value of tie strength to each one of the participant's friends. In a nutshell, the input of BFF is the profile of the participant in the social network, and the output is a set of user groups and a value of tie strength for each one of the participant's contacts.

Figure 1 shows an overview of BFF and how it works. The interface between BFF and the user is a web page. As the figure shows, BFF collects information from the user's Facebook account. We chose Facebook as the first SNS for experimentation and for our first development of BFF due to the success of this SNS and its popularity. Nevertheless, BFF can be easily adapted to other social networks and even to social networks with a distributed architecture, like for example Friendica³. Therefore, before users can use BFF, they have to log in Facebook and give permission to BFF to access their Facebook information. Once the permission is given, BFF requests information from the Facebook server. When all the necessary information has been collected, the information is passed to the community prediction module and to the tie strength prediction module. These modules predict a set of groups and tie strength values for the friends of the user. The predictions are shown to the user as a suggestion using again the web interface. The dotted line represents the possibility for the

³ <http://friendica.com/>




Fig. 1 BFF Overview

Community5

Community's new name

Remove	Name	Tie Strength	New Tie Strength
<input type="checkbox"/>	Salvatore	3	-- ▾
<input type="checkbox"/>	Eddy	1	-- ▾
<input type="checkbox"/>	Vicente	1	-- ▾

Add a friend to this community: -- ▾

Fig. 2 Result Sample

user to refine the suggestions created by the two modules. These modifications are stored in the database for future reference.

Figure 2 shows the screen where the results of BFF are presented to the user. In this example, the figure only depicts one of the communities automatically created by BFF. Part of the name of the members of the community has been hidden to preserve their privacy. As shown in Figure 2, the members of the community are sorted by their tie strength value. Users can refine the

results by changing the name of the community, removing/adding members from/to the community, or changing the tie strength value for any member in that community.

The following sections 3 and 4 explain with more details how the tie strength prediction module and the community prediction module work.

3 Tie Strength Prediction Module

As stated in the introduction, BFF predicts the tie strength of the relationships of the user with each person that is socially connected to her. We model tie strength as a linear combination of predictive variables. These variables are information collected from the profile user. During the creation of BFF the usability of our tool was a key factor. BFF has to be capable of predicting the tie strength accurately in a reasonable amount of time, and every user should be able to get an accurate prediction.

The selected predictive variables for BFF are based on the variables proposed in [8]. In their work, Gilbert and Karahalios propose a set of 74 predictive variables. The authors did not consider the cost of collecting the variables and their generalization, they only considered the predictive capabilities of the variables. We tested the average temporal cost of collecting the variables proposed in [8]. On average, only the collection task took 1,210.73 seconds ($\sigma = 435.55$). This amount of time is not acceptable for an experiment with human participants. The participants may lose interest if the experiment is too long. Therefore, we reduced the complexity of our model taking into account less variables and selecting those that require less effort to collect and to process. We detected that the most relevant factor for the temporal cost was the number of friends. Therefore, in order to reduce the time needed for data collection, we removed variables that are collected from the profiles of the user's friends. For example, we did not collect the publications made by the participant on the walls of her friends. Instead, we only collected the publications of the participant's friends on her wall. We also removed profile dependent variables, such as home city or current job. This type of variables are usually left in blank by SNS users. This can lead to incomplete information and prediction errors [6]. We also took the decision of considering only variables that could be easily processed. More specifically, we limited the variables to those that can be simply counted. This decision provides us two advantages: the variables can be easily generalized and we largely reduce the cost of processing the collected information. For example, taking into account variables that depend on the content of a message require a natural language analyze process. Moreover, this type of variables can limit the different users that are able to use BFF (e.g., only English speakers). Instead, counting the number of messages takes less effort and can be obtained from any user profile. Section 5.3.3 offers the results of an empirical evaluation of the temporal cost of information collection.

BFF collects the information from Facebook using Facebook Query Language (FQL). This language enables developers to use a SQL-style interface to query the data stored in Facebook databases. FQL queries are sent through HTTP requests. The results of FQL are paginated, thus, retrieving the entire set of elements of a specific table (e.g., photos, messages, etc.) can require several queries. Active Facebook users tend to have hundreds of friends, pictures, and hundred of posts on their walls. Therefore, collecting all the available information of the user on Facebook can take several minutes (each HTTP query takes a few seconds). We did not know the amount of information that the participants of the experiment would have on their profiles. Therefore, to avoid an excessive data collection cost, we limited the number of queries to a maximum of 20 pages per item.

Applying the restrictions explained above, we had 12 variables left. These 12 variables were not enough to cover the seven tie strength dimensions defined by Granovetter and related literature [3][28][16]. Specifically, two dimensions were not covered: *social distance* and *emotional support*. To cover social distance we included the variable *educational difference*. We included this variable since in the study of Gilbert [8] this variable appeared the most predictive variable in its dimension. However, as explained in section 5.3.1, in the end, this variable can be removed from our model since it has a very low prediction value. To cover emotional support we chose the variable “likes” which is not considered in [8]. Since we only counted the likes given from participants’ contacts to the items collected for the other variables, this variable did not have an extra collection cost. Table 1 shows and explains the 14 selected predictive variables. Table 2 shows the tie strength dimensions and the predictive variables that belong to each dimension.

The equation below represents the tie strength s_i of the i^{th} friend. R_i stands for the vector of fourteen predictive variables of the i^{th} friend. μ_M is the mean strength of mutual friends between the user and the i^{th} friend. Finally, β is the vector of weights applied to the predictive variables and γ is the weight applied to the mean strength of mutual friends. In order to set the weight of each variable we used the findings of [8] as we wanted to avoid the use of a model that completely lacked information on the relative importance of each variable to predict tie strength.

$$s_i = \beta R_i + \gamma \mu_M$$

$$M = \{s_j : j \text{ and } i \text{ are mutual friends}\}$$

After collecting the predictive variables for the friends of the user, the variables are normalized. Then, the tie strength is calculated for each user. The results are normalized to a numeric scale 1-5, where 1 represents that both persons are very distant (mere acquaintance) and 5 that they are very close. The results are presented graphically, as shown in Figure 2, so that users are sorted by group and by tie strength. It is easier to figure out the value of the tie strength of a person by comparing it to the values of the tie strength of

Variable	Explanation
Last communication	Measures the recency of the communication.
First communication	Is an approximation of the duration of the friendship.
Wall messages	Counts the number of messages exchanged using the wall.
Photos together	Counts the photos where both persons (participant and friend) are tagged.
Links shared	Counts the number web page links traded between the friend and the participant.
Initiated wall posts	Counts the number of publications posted by the friend on the participant's wall.
"Likes"	Counts the number of likes given by the friend to the participant's publications.
Inbox messages exchanged	Counts the number of private messages traded between both persons.
Inbox thread depth	Measures the length of the conversations between both persons.
Number of friends	Is the total number of friends of the friend.
Number of common friends	Counts the number of friends that are common for both persons.
Photo comments	Is the number of comments made by the friend to the photos of the principal user.
Educational difference	Measures the difference in a numeric scale: none = 0, high school = 1, university = 2, PhD = 3.
Mean tie strength of mutual friends	Taking into account the mean tie strength of the friends that are common for both persons we can capture the idea of how relationships are modified by the social cliques.

Table 1 Predictive variables considered.

Dimension	Variables
Intimacy	First communication. Number of friends. Photos together.
Intensity	Wall messages. Initiated wall post. Inbox messages exchanged. Inbox thread depth. Photo comments.
Duration	First communication.
Social distance	Educational difference.
Reciprocal services	Links shared
Emotional support	Likes
Structural dimension	Mean strength of mutual friends. Number of mutual friends.

Table 2 Predictive variables and tie strength dimensions

the relationships with others. As in the grouping step (explained below), the participant can refine the results of the tie strength calculation.

4 Community Prediction Module

The community prediction module is in charge of dividing the network of relationships of the user into communities. This module queries Facebook about the friends of the user and the friends of those friends (mutual friends). With this information the module builds a graph where the nodes are the friends of the user, and the connections between the nodes are friend relationships. This graph is used as the input for the community finding algorithm. The output of the algorithm, a partition of the graph, is shown to the user. The user can modify the communities proposed by the algorithm.

As in the tie strength prediction algorithm, we wanted to present to the participants of the experimental evaluation a suggestion that they can modify. Creating every community from scratch can be a challenging task and we wanted to avoid participants of getting tired of the experiment. The algorithm proposed by Shen et al. in [24] was chosen as the initial community finding algorithm for BFF. The algorithm is founded on the triadic closure principle, which suggests that, in a social network, there is an increased likelihood that two people will become friends if they have friends in common. Based on the results obtained by the authors, this algorithm performs accurately on natural created communities which is the type of communities that the community prediction module has to manage.

According to the results of Shen et al. [24], their algorithm performs better than Infomap [22] and Louvain [1] algorithms. These two community finding algorithms (Infomap and Louvain) are two of the best ones [14]. On the one hand, Infomap uses the probability flow of random walks on a network as a proxy for information flows in the real system and decompose the network into modules by compressing a description of the probability flow. On the other hand, Louvain algorithm is founded on a heuristic method that is based on modularity optimization. The modularity of a partition is a scalar value between -1 and 1 that measures the density of links inside communities as compared to links between communities [9]. As we did not know if the Shen's algorithm was going to be accurate with Facebook communities, we also tested Infomap and Louvain algorithms. The results of the test comparing the three algorithms are shown in the evaluation subsection 5.3.2.

5 Experimental Evaluation

The goal of our experimental study is to evaluate the accuracy of our BFF tool in terms of community and tie strength prediction. Specifically, we want to answer the following questions:

- How accurate is the community module in grouping the contacts of a user?
- How accurate are the predictions of the tie strength module?
- Do users perceive that BFF is a good tool in general? In other words, do they think that BFF is capable of inferring accurate information from their available data on Facebook?

To answer these questions, we performed an experimental evaluation with Facebook users. In the rest of this section, we firstly introduce the experimental settings and, after that, we report the main findings.

5.1 Participants

Our 38 participants were recruited using a Facebook page as well as posters posted on the Universitat Politècnica de València university. We used the viral properties of publications on Facebook to attract participants out of the college environment. Participants were rewarded with a gift voucher for El Corte Ingles (a famous chain of shopping centers in Spain). The participants also entered into one Nexus7 tablet raffle.

The participants filled a form with demographic information. The sample consisted of 9 women (24%) and 29 men (76%). 52% of the participants had an age between 18 and 24, 25% had an age between 25 and 29, 22% had an age between 30 and 39, and 1% of the participants had an age between 40 and 49. Regarding studies degree, the majority of them had a college degree (71%), 14% of them had a PhD, other 14% of participants had high school degree and 1% of the participants had a primary school degree. Finally, 76% of the participants were students and the other 24% were working. We consider that the Facebook viral effect succeeded as at least every age, study degree and professional status was represented. This is specially important in the academic environment, when very often, experimentation with humans tend to be limited to college students. However, attending to the study made by Johnson et al. [12], two demographic groups were underrepresented, the group of women and the age group 40+.

Regarding the number of participants' friends on Facebook, the minimum number of friends was 35 and the maximum was 529. The mean number of friends per participant was 232.19. In total, we analyzed 12803 friend relationships. The majority of participants use Facebook regularly, 84% of them enter Facebook several times per day, the other 16% visit Facebook at least once every few days.

5.2 Method

The participants in our experiment had to try BFF and evaluate its precision. BFF was created to ensure that its use would be easy for anyone. The participants only had to access to the web page of BFF, log in with their Facebook account, and start the application.

After BFF completed its process, the participants were requested to correct any possible errors in tie strength prediction and in user grouping. Users could change the tie strength value of any contact, move users freely from one community to another, and create new communities. These possible corrections were stored in order to evaluate the performance of BFF.

Finally, the participants were requested to answer a short survey to find out their opinion about BFF. The survey was composed of the four following questions:

1. How well did BFF group your friends into communities?
2. How well did BFF predict the tie strength between you and your friends?
3. In general, how accurate do you think BFF is?
4. How accurate do you think BFF is considering it only accesses your information on Facebook? For example, if one of your friends on Facebook is your brother, but you have never interacted with him on Facebook, it is impossible for BFF to accurately predict the tie strength between you and your brother.

Each question was rated on a Likert scale 1-5: 1 = very bad, and 5 = very good. The first and second question addressed specific parts of BFF (the grouping feature and the tie strength prediction respectively). The third and fourth questions were general questions. The intention of the fourth question was to clarify the limitations of BFF to the users. Currently, BFF is limited to the bounds of Facebook; therefore, it only considers the interactions and social connections that occur on Facebook. In future work, we expect to collect information from different sources than Facebook, so BFF will be able to avoid this limitation.

5.3 Results

In this section we analyze the results obtained for both modules, tie strength prediction and community prediction. Apart from reviewing the performance of both modules, we also study how their performance can be enhanced. We also analyze the performance of BFF in terms of complexity and temporal cost.

5.3.1 Tie Strength Prediction Module

With regard to tie strength prediction, the module performed very accurately. It achieved a Root Mean Squared Error (RMSE) of 0.6271 and a Mean Absolute Error of 0.1791 on a discrete scale 1-5, where 1 is the weakest and 5 is the strongest. We chose to discretize⁴ the tie strength in order to facilitate the understanding of the results to the users.

We performed a linear regression analysis to observe what variables were more useful for tie strength prediction and to inspect if the initially chosen coefficients were suitable. The β coefficients for the variables are shown in Table 3.

⁴ The discretization process might have caused a higher prediction error. For example, a user with a tie strength of 3.6 and another with a strength of 4.4 will be both assigned a strength of 4 during the discretization process. As future work, we plan to study the effect of discretization in the prediction error, so that we could achieve a trade-off between the understandability of the results and the error introduced because of the discretization.

Variable	β
Last communication	-1.3487
First communication	1.2603
Mean strength of mutual friends	1.0546
Photos together	0.9529
Likes	0.6223
Number of friends	-0.4903
Inbox thread depth	0.4785
Initiated wall posts	0.3138
Inbox messages exchanged	0.2602
Links shared	0.1282
Wall messages	0.1219
Number of common friends	-0.0338
Educational difference	-0.0164
Photo comments	-0.0028

Table 3 β coefficients for predictive variables after regression

Using a model with the coefficients specified in the Table 3 we obtained a RMSE of 0.614. The small difference between the error of the previous model and the model after the regression process shows that the initial coefficients applied to the variables were appropriate.

Table 3 also shows some variables with very low coefficients. It is interesting to study if it is possible to remove these variables from the model without degrading the predictions significantly. As explained before, each variable requires querying the SNS for the information. This process can be time consuming, specially on very active users. Reducing the amount of information needed for tie strength prediction directly improves the time performance of the module. We performed a multilinear stepwise⁵ regression to observe what variables could be removed from the model maintaining an acceptable error rate. Table 4 shows the predictive variables sorted by β coefficient after the stepwise regression. As it can be seen on the table, the variables with lower β coefficient are also the variables with higher p-value. Therefore, we can create a new model without these variables (*links shared*, *wall messages*, *photo comments*, *common friends*, and *educational difference*) and still maintain an acceptable level of accuracy. This new model, with only 9 variables, has an RMSE = 0.7964. Comparing this error rate to the error rate obtained by the complete model, the increase is not significant. Besides the performance benefit, as explained above, removing education degree is specially beneficial as profile dependent variables, like this one, are usually left in blank by SNS users. This can lead to incomplete information and prediction errors [6].

As stated previously, SNS users struggle to set up privacy settings. If the aim of BFF is to lighten the burden of this task, its suggestions cannot contain a huge number of errors that need correction. Therefore, it is crucial to keep to the minimum the number of corrections that the users need to make to the suggestions presented by BFF. Analyzing the experimental data, we found that

⁵ A sequence of F-tests is used to control the inclusion or exclusion of variables

Variable	β	p-value
Last communication	-1.3527	< 0.001
First communication	1.2684	< 0.001
Mean strength of mutual friends	1.0485	< 0.001
Photos together	0.9608	< 0.001
Likes	0.6817	< 0.001
Inbox thread depth	0.5110	< 0.001
Number of friends	-0.4986	< 0.001
Initiated wall posts	0.3412	< 0.001
Inbox messages exchanged	0.2675	< 0.001
Links shared	0.1836	0.1036
Wall messages	0.1791	0.1023
Photo comments	0.0562	0.62
Number of common friends	-0.0303	0.55
Educational difference	-0.0161	0.5161

Table 4 β coefficients and p-values for predictive variables after stepwise regression

the participants made an average of 42.47 tie strength corrections. Considering that the average number of friends of our participants was 232.19, having to correct only 18.29% of friend relationships can effectively speed up the process of classifying friends before setting privacy policies. It is worth noting that BFF only needs 14 predictive variables, which ensures that it can work fast while maintaining a good accuracy. Moreover, the generalization of the predictive variables allows BFF to be accurate without taking into account the specific characteristics of the user.

5.3.2 Community Prediction Module

In order to evaluate the performance of the community prediction module we used the Normalized Mutual Information (NMI). NMI gives a measure of the quality of the partition obtained by the module comparing that partition to the partition made by the user. NMI is in range $[0,1]$, it equals 1 when both partitions are equal and 0 when both partitions are independent. We use the method proposed by Lancichinetti et al [15] to calculate the NMI. This method defines the NMI between two partitions X and Y as:

$$N(X|Y) = 1 - \frac{1}{2}[H(X|Y)_N + H(Y|X)_N]$$

Where $H(X|Y)_N$ and $H(Y|X)_N$ are the normalized conditional entropy. To calculate this entropy, the algorithm considers the differences between the most similar clusters in both partitions. In other words, the algorithm calculates the total entropy between each cluster in partition X and the most similar cluster in partition Y , and vice versa.

As explained in section 4 the participants were asked to apply any needed correction to the communities created by Shen’s algorithm. Besides this algorithm, we chose two additional algorithms (Infomap and Louvain) to compare the performance of all three. We calculated the NMI between the communities




Fig. 3 NMI for different community finding algorithms

created by each algorithm and the final communities specified by the participants. Figure 3 shows the mean NMI obtained by each algorithm. Infomap is the algorithm with best performance as it achieves a mean NMI of 0.55. Infomap yields more and smaller communities than the other two, this behavior seems to coincide more precisely with the natural partitions made by the participants. Since the users were asked to correct the results obtained by Shen’s algorithm, rather than creating their communities from scratch, some bias might have been introduced. It is likely that some communities created by Shen’s algorithm that were slightly different to the preferences of the participants were considered as correct by the participants. It is possible that if Infomap had been chosen as the initial community finding algorithm, its NMI would have been even better. In this case, the participants would have corrected the suggestions offered by Infomap. Hence, a similar bias as the one introduced in the experiment towards Shen’s algorithm would have been introduced towards Infomap algorithm.

In order to maintain BFF usable it is important to keep the number of corrections that users need to make to BFF suggestions to the minimum. We analyzed the number of changes that participants needed to make to the communities suggested. Figure 4 shows the mean proportional number of changes made to community suggestions. As can be inferred from the average NMI obtained by the three algorithms, Infomap outputs needed less corrections than Shen’s and Louvain’s. In average, participants needed to move 32,92% of their contacts to other communities in order to adapt the partition calculated by Infomap to their preferences.

Another factor to consider when selecting the community algorithm is the execution time. We performed a test to measure the execution time for each algorithm using the evaluation data. Figure 5 shows the average execution time for each algorithm in milliseconds. The fastest algorithm is Louvain; however,




Fig. 4 Mean proportional number of community changes




Fig. 5 Average execution time for each algorithm in milliseconds

every algorithm is very fast and their execution time do not affect the general performance of the tool.

As a conclusion, Infomap algorithm is the best algorithm overall. It achieved the highest mean NMI and its outputs required less corrections than Shen and Louvain algorithms. Infomap is not the fastest one, however, the execution time is not relevant as the three algorithms are very fast with the size of Facebook communities and the execution time differences are not significant. Therefore, we have replaced the Shen algorithm with Infomap algorithm for future versions of BFF.

5.3.3 Performance of BFF

The prediction model of tie strength of BFF is based on [8]. Before the experiment with the participants, we tested the average temporal cost of collecting the entire set of variables⁶ proposed in [8]. For this test we used the profiles of 10 members of our research department. The average time of collection (without considering the cost of processing the collected information) was 1,210.73 seconds ($\sigma = 435.55$); with a maximum of more than 30 minutes for a very active profile (more than 800 friends and daily updates). If BFF aims to be a usable tool, this amount of time is not acceptable. Users have to get the results in a reasonable amount of time. Moreover, since the relationships are dynamic and change over time, users might need to use BFF regularly. Thus, BFF has to be an agile and fast tool.

As explained in section 3, we reduced the complexity of BFF removing variables that were costly to collect or to process. We reduced the number of variables for tie strength prediction to 14. We tested the collection time for these 14 variables using the same 10 profiles of members of our research group. The average collection time was 210,48 seconds ($\sigma = 65.94$). On average, collecting this set of 14 variables was 5.77 ($\sigma = 1.28$) times faster than collecting the 32 variables proposed by Gilbert and Karahalios. As a conclusion, the new selection of variables enables BFF to perform its tasks in a reasonable amount of time (on average 3.5 minutes) while offering accurate results.

5.3.4 Survey Results

The participants also rated the performance of BFF by answering a short survey. Each question was rated using a Likert scale 1-5. Figure 6 shows the mean rating for each survey question. The results reflect that the participants rated BFF performance positively. The participants perceived a slightly better accuracy in tie strength prediction than in community prediction. This perception can improve after the replacement of the community finding algorithm, as the Infomap algorithm has proved to perform better. Another result to note is that the participants rated the second general question (question 4) higher than the first general one (question 3). When answering the first general question, the participants did not consider the limitations of BFF. Therefore, even when almost every friend was rated correctly, they detected mistakes. Due to the brief explanation in the second general question about how BFF works, the participants realized that BFF is limited by the bounds of Facebook, and, for example, that it cannot predict the tie strength of a relationship that mainly occurs outside Facebook. When the participants became aware of the limitations of BFF, they took into account how they interacted with others. They achieved an accuracy of 86% on Facebook in order to make their judgments. This explains the better rating for the second general question.

⁶ Gilbert and Karahalios state in their paper that they consider 74 variables; however, in the paper they only show and explain 32 of these variables. In the end, we tested the information collection time considering these 32 variables.




Fig. 6 Score for the survey questions

5.4 Generalizability and Limitations

Since Facebook is no longer limited to a specific group of users, it is necessary to take into account the participants of other demographic groups. Despite the fact that the group of women and people 40 years old or older were underrepresented, we believe that our sample accurately reflects the current demographics of Facebook [12].

Another element that increases the generalizability of our study is that it is not based on surveys. The utility of surveys is unquestionable; however, the results obtained through surveys can introduce bias. The data analyzed in this paper is real data retrieved from real Facebook users. We collected a large volume of information that is not possible to obtain through surveys (thousands of relationships and their characteristics).

It is possible that our method introduced bias. Participants of the experiment were asked to correct the results inferred by BFF, thus, their perception of their relationships could be affected by the results shown. The tie strength values specified by the users could be modified by the results shown. Also, as explained above, the users were asked to correct the results obtained by Shen's algorithm, rather than creating their communities from scratch. Despite the possible bias introduced, the aim of BFF is providing good enough recommendations that help users to classify their relationships with others on a SNS. Therefore, the measure of the performance of BFF has to be based on how appropriate the recommendations of BFF are for users and the number of corrections that the users need to perform to these recommendations.

6 Related Work

Recent works have proposed models to predict tie strength. Gilbert et al.[8] proposed a model, based on Granovetter’s work, that predicted tie strength among the users of Facebook. The authors identified a set of 74 predictive variables that can be found on Facebook. They performed an experiment to infer the tie strength of a portion of the participants’ friends. They achieved an accuracy of 86%. Another work that predicts tie strength of social links is [13]. Like in the work of Gilbert, the authors define a set of 50 predictive variables. In this work the authors aim to discriminate strong links from weak links. However, they do not consider a scale in the strength of the link, they are either strong or weak. These two works use a supervised learning model that needs human intervention to work properly. Aiming at the same objective, Xi-ang et al.[30] proposed a model to infer relationship strength based on profile similarity and interaction activity, with the goal of automatically distinguishing strong relationships from weak ones. It is worth noting that this model relies on an unsupervised learning method, but it lacks a empirical evaluation with real users. All three works show that it is possible to infer tie strength from the available personal data in a SNS. These three works differ from ours in that they aim to create models to predict tie strength from the information available on a SNS. However, they do not offer tools that social network users can use to help them to form friend groups and set privacy policies. Moreover, they only consider the predictive capabilities of the variables chosen for their models, but they do not take into account factors like the computational cost of collecting these variables, which is an important factor when creating an usable tool.

The other main feature of BFF is that it suggests friend groups to the participant user. The main idea is that with the grouping and tie strength information the user has enough elements to create appropriate privacy policies. The work of Fang and coworkers [6] proposes a tool that suggests privacy policies for certain elements of a Facebook user profile. This work bases the privacy suggestions in grouping user’s contacts in contexts. Every contact in the same context is granted the same access permissions. The authors present a tool called Privacy Wizard that helps user to set the privacy policies to protect user’s traits, like birth date, address, and telephone number. This work does not consider tie strength, and as the authors proved in [29], it is a key variable to consider when determining the disclosure degree of the elements being shared in a social network.

Other works present mechanisms that can partially infer users’ social network and its characteristics from sources of information different than SNSs. In [4] the authors propose a method that extracts a social network for a user given her mailbox and the information available on Internet. A similar approach is presented in [19]. In this work the authors present Polyphonet. From a given set of persons, the authors find the social connections among them by querying to Google. The authors estimate the strength of the relationship between two persons by co-occurrences of their two names. These two works differ from

ours in that they do not rely in a SNS to extract social information from users. However, this approach also has limitations. Relying on information sources that do not necessarily contain social relevant information may lead to errors. For example, two persons may appear in several web pages together but do not have any social link. In order to avoid this problem, both works ([19,4]) require a predefined set of persons that will form the social community. In contrast, relying only on Facebook data guarantees that the social links will actually exist, but may also lead to errors. Even when the connection truly exists, the interactions between two persons may occur outside Facebook. Therefore, the strength of such link will be incorrectly predicted by our software. In the future, we plan to expand the search of variables for defining the groups and the tie strength with information that can be found outside the social network, like the information available in the participant's mailbox or in the personal web page of a user of the social network.

Virtual worlds, such as Second Life, are social networking sites that offer different ways of communication in a virtual environment. People can explore, participate in social activities, create and trade properties and services. Sharma et al. [23] explore business possibilities in this environment. The researchers found that different ways of communication in this virtual world have a direct correlation with the confidence between potential customers and companies. In the same way BFF is focused on facebook-like SNSs, a new tool could be developed to infer the tie strength and friend communities of users of virtual worlds.

The work of Murukannaiah and Singh [20] presents Platys Social. The authors developed a software that runs on a mobile device. This software learns a user's social circles and the priority of the user's social connections from daily interactions. The software infers the interactions from information that is available on mobile devices, such as wi-fi networks, bluetooth connections, phone calls, and text messages. The work of Murukannaiah and Singh presents a new approach for extracting social information from the real world, and not only from Internet. Their work and ours could be merged so that tie strength could be computed taking into account day by day encounter frequency and the information stored on a SNS like Facebook.

7 Conclusions and Future Work

In this paper, we have presented the BFF tool that is able to automate the elicitation of tie strength and user communities to a large extent. In particular, we evaluated it using real-world data from real users of Facebook. BFF achieved 81.71% accuracy in tie strength (i.e. users only needed to correct 18.29% of their relationships). Moreover, the tie strength prediction model used by BFF was composed of only 14 variables, and it could be even simplified to use 9 variables with a slight loss of precision.

Regarding community prediction, BFF achieved 67.08% accuracy (i.e. users only needed to move 32.92% of their contacts to different communities). We

Feature	Accuracy
Tie strength prediction with 14 variables	81.71%
Community prediction	67.08%

Table 5 Results overview.

evaluated three different community finding algorithms (Infomap, Louvain and Shen) in order to find the best one for the objective of BFF. Infomap algorithm outperformed the other two algorithms achieving a better accuracy. Finally, according to a survey performed by the participants in the experiments, we obtained that users consider BFF as providing good predictions of tie strength and communities. Table 5 shows an overview of the results.

Many research paths open from here. The first one, and the motivation of this work, is to develop intelligent personal agents that will aid users in the definition of their privacy policies for SNSs. To this aim, this intelligent agent will use the extracted information to recommend adequate privacy policies. Another path for further research is to improve the predictive capabilities of BFF by collecting information from other sources than the SNS (such as users' mailbox, personal web pages, Internet search engines and mobile devices [4, 19, 20]).

Finally, it is worth noting that the information that our tool provides can also be used in many other agent-based applications. For instance, the tie strength among agents is needed to obtain the optimal social trust path in complex social networks [18]. Moreover, in automated negotiation environments, agents could judge the outcome of a negotiation as being distributively fair based on the tie strength between them [25]. Apart from being used in agent-based applications, BFF could also be used in many other more general applications. For instance, it can be very useful to perform experiments with humans in which either tie strength or user communities are needed to evaluate the results of the experiments (such as in [29]). In this case, BFF can speedup the experiments by automating part of the process for eliciting tie strength and user communities from the participants.

8 Acknowledgements

This work has been partially supported by CONSOLIDER-INGENIO 2010 under grant CSD2007-00022, and TIN 2008-04446 and PROMETEO/2008/051 projects. Ricard L. Fogués was awarded a 4-year PhD studentship (FPI) by Universitat Politècnica de València under Programa de Ayudas de Investigación y Desarrollo (PAID).

References

1. V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*,

- 2008(10):P10008, 2008.
2. D. Boyd and E. Hargittai. Facebook privacy settings: Who cares? *First Monday*, 15(8), 2010.
 3. R. Burt. *Structural holes: The social structure of competition*. Harvard Univ Pr, 1995.
 4. A. Culotta, R. Bekkerman, and A. McCallum. Extracting social networks and contact information from email and the web. 2004.
 5. N. Ellison, C. Steinfield, and C. Lampe. The benefits of facebook friends: social capital and college students use of online social network sites. *Journal of Computer-Mediated Communication*, 12(4):1143–1168, 2007.
 6. L. Fang and K. LeFevre. Privacy wizards for social networking sites. In *Proceedings of the 19th international conference on World wide web*, pages 351–360. ACM, 2010.
 7. S. Fortunato. Community detection in graphs. *Physics Reports*, 486(3-5):75–174, 2010.
 8. E. Gilbert and K. Karahalios. Predicting tie strength with social media. In *Proceedings of the 27th international conference on Human factors in computing systems*, pages 211–220. ACM, 2009.
 9. M. Girvan and M. Newman. Community structure in social and biological networks. *Proceedings of the National Academy of Sciences*, 99(12):7821, 2002.
 10. M. Granovetter. The strength of weak ties. *American journal of sociology*, pages 1360–1380, 1973.
 11. R. Gross and A. Acquisti. Information revelation and privacy in online social networks. In *Proceedings of the 2005 ACM workshop on Privacy in the electronic society*, pages 71–80. ACM, 2005.
 12. M. Johnson, S. Egelman, and S. Bellovin. Facebook and privacy: it’s complicated. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, page 9. ACM, 2012.
 13. I. Kahanda and J. Neville. Using transactional information to predict link strength in online social networks. In *Proceedings of the Third International Conference on Weblogs and Social Media (ICWSM)*, 2009.
 14. A. Lancichinetti and S. Fortunato. Community detection algorithms: A comparative analysis. *Phys. Rev. E*, 80:056117, Nov 2009.
 15. A. Lancichinetti, S. Fortunato, and J. Kertsz. Detecting the overlapping and hierarchical community structure in complex networks. *New Journal of Physics*, 11(3):033015, 2009.
 16. N. Lin, W. Ensel, and J. Vaughn. Social resources and strength of ties: Structural factors in occupational status attainment. *American sociological review*, pages 393–405, 1981.
 17. H. Lipford, A. Besmer, and J. Watson. Understanding privacy settings in facebook with an audience view. In *Proceedings of the 1st Conference on Usability, Psychology, and Security*, pages 1–8. USENIX Association Berkeley, CA, USA, 2008.
 18. G. Liu, Y. Wang, and M. Orgun. Optimal social trust path selection in complex social networks. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence, AAAI*, pages 1391–1398, 2010.
 19. Y. Matsuo, J. Mori, M. Hamasaki, T. Nishimura, H. Takeda, K. Hasida, and M. Ishizuka. Polyphonet: An advanced social network extraction system from the web. *Web Semantics: Science, Services and Agents on the World Wide Web*, 5(4):262 – 278, 2007. World Wide Web Conference 2006 Semantic Web Track.
 20. P. Murukannaiah and M. Singh. Platys social: Relating shared places and private social circles. *Internet Computing, IEEE*, (99):1–1, 2011.
 21. D. Quercia, R. Lambiotte, M. Kosinski, D. Stillwell, and J. Crowcroft. The personality of popular facebook users. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work (CSCW’12)*, 2012.
 22. M. Rosvall and C. Bergstrom. Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences*, 105(4):1118–1123, 2008.
 23. G. Sharma, Y. Qiang, S. Wenjun, and L. Qi. Communication in virtual world: Second life and business opportunities. *Information Systems Frontiers*, 15(4):677–694, 2013.
 24. K. Shen, L. Song, X. Yang, and W. Zhang. A hierarchical diffusion algorithm for community detection in social networks. In *Cyber-Enabled Distributed Computing and Knowledge Discovery (CyberC), 2010 International Conference on*, pages 276–283. IEEE, 2010.

25. C. Sierra and J. Debenham. The LOGIC negotiation model. In *AAMAS '07: Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, pages 1–8. ACM, 2007.
26. J. Staddon, D. Huffaker, L. Brown, and A. Sedley. Are privacy concerns a turn-off?: engagement and privacy in social networks. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, page 10. ACM, 2012.
27. K. Strater and H. R. Lipford. Strategies and struggles with privacy in an online social networking community. In *Proceedings of the 22nd British HCI Group Annual Conference on People and Computers: Culture, Creativity, Interaction - Volume 1*, BCS-HCI '08, pages 111–119, Swinton, UK, UK, 2008. British Computer Society.
28. B. Wellman and S. Wortley. Different strokes from different folks: Community ties and social support. *American journal of Sociology*, pages 558–588, 1990.
29. J. Wiese, P. Kelley, L. Cranor, L. Dabbish, J. Hong, and J. Zimmerman. Are you close with me? are you nearby? investigating social groups, closeness, and willingness to share. In *Proceedings of the 13th international conference on Ubiquitous computing*, pages 197–206. ACM, 2011.
30. R. Xiang, J. Neville, and M. Rogati. Modeling relationship strength in online social networks. In *Proceedings of the 19th international conference on World wide web*, pages 981–990. ACM, 2010.