

# What Makes A Group Fail: Modeling Social Group Behavior in Event-Based Social Networks

Xiang Liu  
NYU Tandon School of Engineering  
Brooklyn, NY  
xl493@nyu.edu

Torsten Suel  
NYU Tandon School of Engineering  
Brooklyn, NY  
torsten.suel@nyu.edu

**Abstract**—Event-based online social networks, which are used to maintain interest-based groups and to distribute and organize offline events, have recently gained increasing popularity. In event-based social networks, some groups survive and thrive, while other groups fail. How to build successful groups and what factors make a “healthy” group are important open problems. We address the problem of modeling social group behavior and present detailed studies on group failure prediction by analyzing a large online event-based social network. We investigate both the statistical properties and the structural features of the social groups, and find that event features play an important role in distinguishing social groups with different topics and categories. We also observe that tightly knit communities have less average event participation, and both low level diversity and high level diversity in members’ event participation will harm group activity participation. We then analyze the data of thousands of social groups collected from the Meetup platform with the goal of understanding what makes a group fail. We use two different feature selection methods in this paper and build a model to predict which groups will fail over a period of time. The experimental results show that social group failures can be predicted with high accuracy, and that member features contribute significantly to the success of social groups.

**Keywords**-Social Networks; Information Diffusion; Group Evolution; Event-Based Social Networks

## I. INTRODUCTION

Recently, event-based online social networks (EBSN), such as Meetup ([www.meetup.com](http://www.meetup.com)), have become more and more prevalent. Such services are defined as social platforms that allow people to create interest-based groups to organize and promote offline events ranging from after-work parties to technical conferences. On Meetup, users are allowed to form and join social groups to interact with other users. A social event can be created by specifying a group, event description, location, and time, which is announced to the group members or other users who may be interested in it. Then the targeted users may send RSVPs (“yes”, “no”, or “maybe”) to announce whether they want to go to the event [1]. Figure 1 depicts the social objects in the Meetup social networks. There are over 10,000 groups, 800,000 users, 630,000 past events, and around 2,000 upcoming events in New York City as of October, 2014. As such event-based

social networks grow so rapidly, it is important to study the social group behavior and model how it evolves.

One principal task of studying social networks is to identify communities, which allows us to discover related social objects with similar behaviors, interests, or background. For example, communities in the co-authorship graph may correspond to people who work on the same research discipline, while communities in social networks may correspond to people who are from the same school or hometown, and communities in the web graph may correspond to pages on similar topics. In this paper, we are interested in clustering social groups that share common properties or attributes. Community detection in networks has been extensively studied; a lot of research has focused on detecting social communities using social object interactions [2]. Other methods have focused on detecting communities based on features of social objects [3]. On Meetup, there is no friendship connection between users; instead two users are loosely connected by either the same group affiliation or event co-participation. In addition, there is no direct interaction between two groups except for common group members; thus, we adopt clustering algorithms based on attributes of social groups to detect related social groups in EBSNs. While considering both the statistical properties and the structural features (e.g., average clustering coefficient, entropy of event degree distribution in event-user bipartite graph, etc.) of the social groups, we find that event features play an important role in detecting different clusters of social groups. Social groups that are clustered by event features are more cohesive in terms of both topic distribution and category distribution. We also find the following:

- Tightly knit communities have more frequent events, but fewer average event RSVPs.
- Both low-level and high-level diversity in members’ event participation will harm group event activity participation.
- Event frequency has a strong correlation to the level of diversity in the group’s event participation. When a group has more frequent events, the diversity in event participation within a group increases.

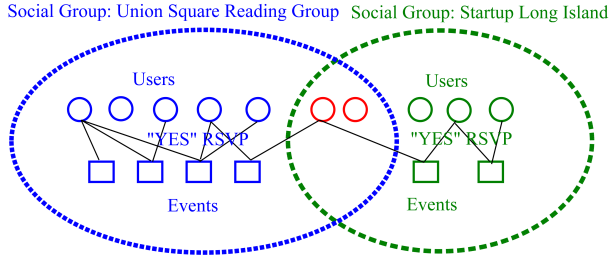


Figure 1. Meetup as an example of an event-based social network. Users and events can be modeled as nodes in a bipartite graph where edges represent positive RSVPs. Shown in the figure are two social groups: Union Square Reading Group and Startup Long Island. For Union Square Reading Group, there are seven users and four events; for Startup Long Island, there are five users and two events. The users marked as red circles are the intersection members of both social groups. Each edge connecting a user and an event represents a positive RSVP.

Another key research problem for social groups is the analysis of group evolution. Most existing work models social group evolution as an information diffusion process and measures group growth from the membership growth point of view [4]. However, on Meetup, users can join as many groups as they want; therefore, group sizes can increase monotonically. Hence, the group size does not necessarily reveal how successful the group is, if it does not attract users to participate in events. Analyzing the data of thousands of social groups collected from the Meetup platform in New York City with the goal of understanding what makes a rising or stable group, we first justify a definition of group failure as no events being organized or no event RSVPs being collected for twelve months. Then, we investigate the question of how some groups survive, while other groups fail, over a period of time. What are the key factors that predict whether a group will succeed or fail? Inspired by [5], we use a novel model to understand the relationship between group structural features and its future growth. Then we build a classifier to predict which groups will fail over a period of time. The experimental results show that our model can predict social group failures with decent accuracy.

The rest of this paper is organized as follows. In Section 2 we discuss related work. We explain our dataset and feature selection in Section 3. Then we describe our social group clustering algorithms and results in detail in Section 4. In Section 5, we present our group failure prediction model and the experimental results. Finally, we provide concluding remarks in Section 6.

## II. RELATED WORK

Given a social network community, many studies have tried to predict its evolution in terms of membership dynamics. For example, [6] utilized decision tree-based methods to analyze how groups evolve over time, using the DBLP dataset. [5] designed and ranked features that can predict fu-

ture group stability. [7] tried to identify factors contributing to the growth and longevity of groups within social networks, using NING community data. [8] presented network analyses of information diffusion on Twitter capturing three major properties: speed, scale, and range.

The existing social theory on group success and failures has also pointed out that a malfunctioning group can be caused by too many group members, or, if the group leader fails to enforce a common purpose [18].

Most of the existing work models community evolution as a diffusion process, where friendship ties across the community boundary attract new individuals to join in the community [9]. However, for EBSNs, e.g., Meetup, the group growth cannot be simply modeled as increase in group membership and event participation, as there are two types of social interactions: online interaction (social group membership) and offline interaction (group event participation) [1]. Also, most existing work models the group evolution based on the DBLP dataset, where conferences are considered as communities, which has much lower dynamics compared with real social networks such as Facebook, Twitter, or Meetup.

Compared to previous studies, we investigate the group evolution problem on a more complex and dynamic social network with more types of social objects. We build classifiers based on social group statistical features, as well as bipartite graph structural features, to predict whether a group will fail (i.e., have no events or event RSVPs for 12 months) in the future.

## III. ANALYSIS OF MEETUP GROUPS

We study social group clustering and social group failure prediction problems in a large-scale EBSN: the Meetup NYC dataset. In the following, we first introduce our dataset, and then propose our community detection methods.

### A. Dataset

Meetup is the world's largest EBSN today, having more than 170,000 online groups with more than 500,000 monthly offline events. [1] crawled all the Meetup data from Oct 2011 to Jan 2012. However, the published dataset only contains user-event pairs, user-group pairs, user/group-tag pairs, and user/event locations, while all properties of social groups (e.g., group name, group description, when the group was created, etc.) are missing. Thus we take a similar data allocation strategy as stated in [10] to collect Meetup data (groups, events, and members) with node metadata (group/event/member properties, event RSVPs, and member topics) in New York City, starting from 2002 to October 2014. We extracted all groups of New York City. Then starting from the crawled groups, we extracted all the users and events associated with each social group. Table I shows the statistics of the dataset.

We then filter out “empty” groups that have no events, “empty” events with no RSVPs, and “inactive” users who never participate in events. Additionally, we adopt the concept of an h-index in scholarly impact analysis [11], which measures both productivity and citation impact of the published work of a scientist. Similarly, a social group has an index  $h$  if  $h$  of its  $N_e$  events have at least  $h$  positive RSVPs, and the other  $(N_e - h)$  events have no more than  $h$  positive RSVPs each.

Table I  
STATISTICS OF THE DATASET

Statistic	Meetup NYC
Number of Groups	10,942
Number of Distinct Users	782,995
Number of Events	627,416
Number of RSVPs	7,921,154
Number of Groups with Events	7,643

## B. Social Group Features

Inspired by [5], in order to explore the community structures of social groups in Meetup, we consider a variety of features and categorize them into four types: group features, event features, member features, and structural features.

1) *Group Features*: Group features reveal the basic statistical properties of a social group. All the features are efficient to compute. We mainly focus on the following:

- Number of group members: This is the size of a social group.
- Number of “active” group members: This feature reflects the number of “real” members, who have participated in at least one event, in a social group.
- Number of events: This feature measures the activity level of a social group.
- Length of existence: This feature represents the number of days since a social group was created.
- Group join mode: This is the join mode of a social group, which can be “open”, “closed”, or “approval”.
- Group visibility: This is the visibility of a social group, which can be “public”, “public limited”, or “members only”.

2) *Group Member Features*: Member features reflect the properties of all members in a social group. In this paper, we mainly focus on the following:

- Average event attendance of group members: This measures the average activity level of all members in a social group.
- Standard deviation of event attendance of group members: This features measures how consistent the members’ activity levels are across a social group.
- Average join time of group members: This is the average absolute join time of all members in a social group.

- Standard deviation of member join times: This measures the variation in members’ join time of a social group.

3) *Group Event Features*: Event features reflect the diversity or homogeneity of activities in a social group. We mainly investigate the following:

- Average event positive RSVPs: This is the average number of participants for the events in a social group.
- Standard deviation of event positive RSVPs: This measures how consistent the events’ involvement levels are.
- Average event capacity: This is the average number of positive RSVPs that the events can have before members will be added to the waiting list.
- Standard deviation of event capacity: This measures the variation in the number of positive RSVPs the events can have.
- Average event duration: This is the average event duration in seconds for all the events.
- Standard deviation of event duration: This measures the variation in event duration for all the events.
- Average time between two consecutive events: This feature calculates  $\#events/group$  length, which measures how frequent the events are in a social group.

4) *Structural Features*: Finally, we investigate the structural features of a social group. First, we construct a bipartite graph  $G = (U, V, E)$  as shown in Figure 1, where  $U$  is the event set,  $V$  is the user set, and  $E$  is the set of event participation connections (RSVPs) between events and users. In the event-user bipartite graph, event degree is also known as the event size, while member degree represents how many events a member has attended.

- Entropy of event degree distribution: This feature calculates the entropy of the event participation distribution for a given social group.
- Entropy of member degree distribution: This feature calculates the entropy of the member involvement distribution for a given social group.
- Average clustering coefficient: Except for the event-user bipartite graph, we also construct an event co-attendance graph  $G = (U, E)$  for all the members in a social group, where  $U$  is the user set, and  $E$  is the set of event co-participation connections between users. The edge weight between two users is assigned as the number of events co-attended by two of them. We then measure the clustering coefficient at each member node and then calculate the average of the clustering coefficients across all social group members. The clustering coefficient at each user node can be used to quantify how close its neighbors are to being a clique [5], [12].
- Entropy of user-event participation matrix: This feature measures the entropy of the event participation distribution for the given social group.

#### IV. CLUSTERING SOCIAL GROUPS

In this section, we investigate the community properties of social groups in EBSNs. Because a social group is already a community, here we use the term “community” to represent a set of social groups that share common properties or attributes. As discussed in the previous section, each social group is represented by points in a  $d$ -dimensional vector space, where  $d$  is the feature dimension. Thus, we apply the standard K-means [13] algorithm to cluster all the social groups in New York City with  $h$ -index greater than or equal to five. Each cluster is a set of  $d$ -dimensional vectors,  $D = \{x_i | i = 1, \dots, n\}$ , where  $x_i \in \mathbf{R}^d$  denotes the  $i$ -th social group. The K-means algorithm clusters all the data points in  $D$  such that each data point falls in one of  $k$  clusters.

##### A. Community Detection Evaluations

1) *Davies-Bouldin (DB) Index*: To evaluate the community detection performance, we use the standard Davies-Bouldin (DB) index as defined in [1], [14], which is used to measure the cohesiveness in communities. The formula of the DB index is as follows:

$$DB = \frac{1}{k} \sum_{i=1}^k \max_{j \neq i} \{D_{i,j}\}$$

where


$$D_{i,j} = \frac{\bar{d}_i + \bar{d}_j}{d_{ij}}$$

and  $\bar{d}_i$  is the average distance between each point in the  $i$ th cluster and the centroid of the  $i$ th cluster. Additionally,  $d_{ij}$  is the Euclidean distance between the centroids of the  $i$ th and  $j$ th clusters. The maximum value of  $D_{i,j}$  represents the worst-case within-to-between cluster ratio for cluster  $i$ . The optimal clustering solution has the smallest DB index value, which indicates a more cohesive community.

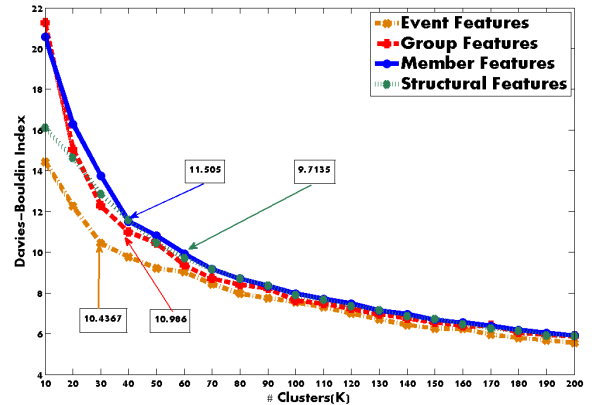
2) *Ground Truth Labels*: We use two sets of ground truth labels to measure the quality of social group communities. First, we collect the groups’ provided topics as the latent community semantics. There were 16,795 unique topics collected from the 2,983 Meetup groups with  $h$ -index greater than or equal to five. Thus, each Meetup group is treated as a  $g_i = \{t_{i1}, t_{i2}, \dots, t_{in}\}$  where  $t_{ij} = 1$  if group  $g_i$  is tagged with topic  $t_j$ .

Besides topic tags, we also use categories as another external ground truth label. There are 25 pre-defined categories for Meetup groups. Similar to topics, each Meetup group is treated as a  $g_i = \{c_{i1}, c_{i2}, \dots, c_{im}\}$  where  $c_{ij} = 1$  if group  $g_i$  is in category  $c_j$ .

After normalization, the similarity score between two Meetup groups  $g_i, g_j$  is calculated by their cosine similarity.



(a) Community detection performance as evaluated by category distribution.



(b) Community detection performance as evaluated by topic distribution.

Figure 2. Community detection performance. The score inside the rectangle is the DB index under the optimal K based on the “knee” method.

##### B. Results

As shown in Figure 2, the event features achieve the best social groups clustering performance under two evaluation settings: the clusters’ cohesiveness evaluation under topic ground truth labels and under category ground truth labels. The DB index of communities clustered using event features is much lower than using the group-level features or the member features, while the DB index of communities clustered using structural features is only slightly higher than using event features. We can thus draw the conclusions as follows: (1) Social groups behave differently and rely more on their event features and structural features than group features and member features. (2) The community detection performance is better, when evaluated using topic ground truth labels, than category ground truth labels. This is probably due to the sparsity in the group-category matrix, since each group is associated with only one category, while each group can be associated with multiple topics.

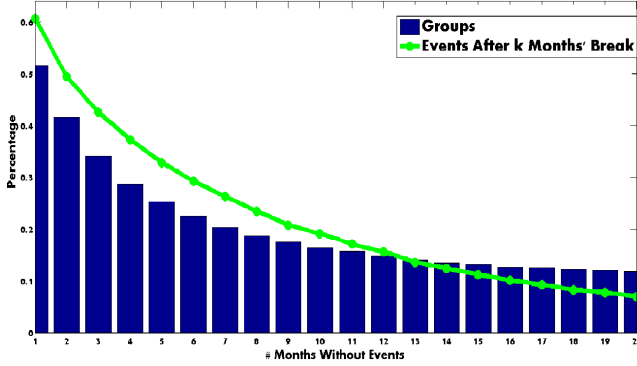


Figure 3. How the number of social groups that have had no event or event RSVP for  $k$ -months changes; how the number of social groups that have had new events again after a  $k$ -months break changes.

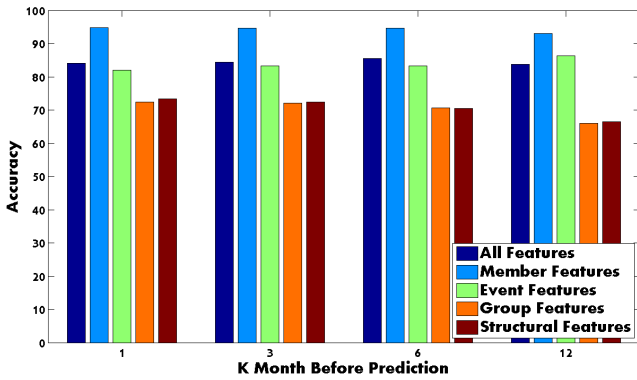


Figure 4. How five categories of features perform in predicting group failure. From the result, we can conclude that the group failure can be reasonably well predicted even 12 months ahead.

## V. SOCIAL GROUP FAILURE PREDICTION

In this section, we introduce our group failure prediction model.

### A. Group Failure Definition

As mentioned earlier, we define group failure or death as “no events or event RSVPs within 12 months.” The validation of this definition can be seen in Figure 3. As shown in this figure, 16% of the 10,942 social groups have had such 12-month break. Meanwhile, only 18% of such “failing groups will have new events again, which makes it a reasonable approach for group “death or “failure.

### B. Predicting Group Failure

We use supervised learning techniques and apply them to feature sets collected one month, three months, and six months prior to possible group failure, respectively. In this paper, we use a random forest classification algorithm [19] to predict group failure. Random forest is a state-of-the-art ensemble learning method for both regression and classification. Because of its robustness to overfitting, it is often used in big data problems [20].

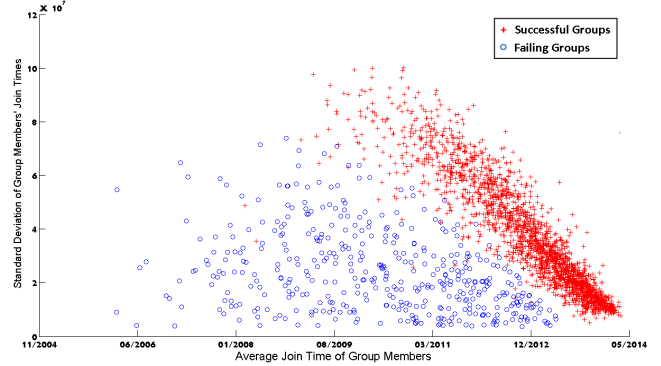


Figure 5. According to our experiment results, member features play an important role in predicting “failing” social groups. By analysing the correlation between different member features and the ground truth label, we find that group success and failure can be easily distinguished by two features: the average join time of group members, and the standard deviation of group members’ join times. The “younger” groups tend to be more “healthy”, while the “old” groups that constantly attract new group members also tend to survive.

1) *Training Data*: Since the number of failing social groups is much smaller than “healthy” groups, we create a training dataset with equal numbers of positive and negative failure groups, following the methodology in [16], [17], so that random prediction attains 50% prediction accuracy. We also create the testing dataset in the same way, with balanced positive and negative samples.

2) *Feature Selection*: We use the same features as in Section 3, which are: event features, member features, group features, and structural features.

In order to understand which features contribute more in modeling group “success” or “failure”, the correlation coefficients between each feature and the class labels (1 for fail, 0 for succeed) are shown in Figure II. We can see that member features and event features are highly significant.

Table II  
TOP FEATURES RANKED BY HOW SIGNIFICANT THE CORRELATIONS WITH THE CLASS LABEL ARE

Rank	Feature	Category
1	Average join time of members	Member
2	Average time between two events	Event
3	Entropy of user-event matrix	Structural
4	SD of member joined times	Member
5	Event Number	Event

### C. Evaluation

We have described the models and the features we use. In this subsection, we describe our experiments for evaluating the performance of the proposed methods.

1) *Evaluation Metrics*: In order to test the effectiveness of the social group failure prediction model, we use the following metric to quantitatively evaluate the results.

$$accuracy = \frac{\#tp + \#tn}{\#tp + \#fp + \#tn + \#fn} \quad (1)$$

2) *Analysis*: We compute group features along with class labels (1 for failure, 0 for rising or stable) from the created date of the groups until one/three/six/twelve months before the prediction time.

We first use the wrapper [21] feature selection approach by utilizing the learning model as a black box to score subsets of features. As Figure 4 shows, we get different results from clustering social groups: member features outperform the other categories in predicting group failure, while event features perform second best. The group features and structural features achieve similar performance. As we can see, group failure can be well predicted even twelve months before the prediction time.

We then use the filter [22] feature selection approach, which is a pre-processing step before the learning step. After selecting the top-10 features using the greedy search algorithm, we obtain slightly better prediction accuracy than by using all features, but the performance is still worse than just group member features. Thus we analyze the correlation between different member features and the ground truth label, and we find that group success and failure is highly correlated with two features, as shown in 5: the average join time of group members, and the standard deviation of group members' join times. It shows that groups with most newly joined members are "healthy", while "old" groups that constantly attract new group members also tend to survive.

## VI. CONCLUSION

In this paper, we have modeled the behavior of social groups in EBSNs using both statistical properties and structural features. We address the problems of community detection and group failure prediction by analyzing a large online EBSN, Meetup. We find that both event-level features and structural features play an important role in distinguishing social groups with different topics or categories. Our research reveals that tightly knit communities have less average event participation but more frequent events, and both low level diversity and high level diversity in members' event participation will harm a group's activity level. We built a classifier based on event features, as well as bipartite graph structural features, to predict whether a group will fail (having no events or event RSVPs during 12 months) in the future. The experimental results have shown that social group failures can be predicted with high accuracy. We notice that the features that work best in predicting group failure are different from those that contribute more to detect communities of social groups. Compared to previous studies, we investigate the group evolution problem on a more complex and dynamic social network with more types of social objects.

## REFERENCES

- [1] X. Liu, Q. He, Y. Tian, W. Lee, J. McPherson, and J. Han. Event-based social networks: linking the online and offline

- social worlds. In ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD 2012.
- [2] A. Vakali, and K. Kafetsios. Emotion aware clustering analysis as a tool for Web 2.0 communities detection: Implications for curriculum development. In World Wide Web Conference, WWW 2012.
- [3] J. Yang, J. McAuley, and J. Leskovec. Community detection in Networks with node attributes. In IEEE International Conference on Data Mining, ICDM 2013.
- [4] E. Zheleva, H. Sharara, and L. Getoor. Co-evolution of social and affiliation networks. In ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD 2009.
- [5] A. Patil, J. Liu, and J. Gao. Predicting group stability in online social networks. In World Wide Web Conference, WWW 2013.
- [6] L. Backstrom, D. Huttenlocher, J. Kleinberg, and X. Lan. Group formation in large social networks: memberships, growth, and evolution. In ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD 2006.
- [7] S. Kairam, D. Wang, and J. Leskovec. The life and death of online groups: predicting group growth and longevity. In ACM WSDM Conference on Web Search and Data Mining, WSDM 2012.
- [8] J. Yang, and S. Counts. Predicting the speed, scale, and range of information diffusion in twitter. In International AAAI Conference on Weblogs and Social Media, ICWSM 2010.
- [9] D. Centola, and M. Macy. Complex Contagions and the Weakness of Long Ties. In American Journal of Sociology, 113(3):702-734, Nov. 2007.
- [10] W. Zhang, J. Wang, and W. Feng. Combining latent factor model with location features for event-based group recommendation. In ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD 2013.
- [11] J. E. Hirsch. An index to quantify an individual's scientific research output. In PNAS 102 (46):16569-16572, 2005.
- [12] D. Watts, and S. Strogatz. Collective dynamics of "small-world" networks. In Nature, 393 (6684):440-442, 1998.
- [13] S. P. Lloyd. Least squares quantization in PCM. In IEEE Transactions on Information Theory 28 (2):129-137, 1982.
- [14] D. Davies, and D. Bouldin. A clustering separation measure. In Pattern Analysis and Machine Intelligence, IEEE Transactions on, 1979.
- [15] A. K. Jain, and R. C. Dubes. Algorithms for clustering data. In Prentice Hall, 1988.
- [16] R. V. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of trust and distrust. In World Wide Web Conference, WWW 2004.
- [17] J. Leskovec, D. Huttenlocher, and J. Kleinberg. Predicting Positive and Negative Links in Online Social Networks. In World Wide Web Conference, WWW 2010.
- [18] M. Sherif, and Carolyn W. Sherif. An Outline of Social Psychology rev.ed. Harper & Brothers: New York pp. 143-180.
- [19] L. Breiman. Random Forests. Mach. Learn., 45(1):5-32, October 2001.
- [20] G. Theodorou, P. S. Thomas, M. Ghavamzadeh. Personalized Ad Recommendation Systems for Life-Time Value Optimization with Guarantees. In Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI-15), 2015.
- [21] R. Kohavi, G. H. John. Wrappers for Feature Selection. Artificial Intelligence, 97 (1-2):273-324, 1997.
- [22] I. Guyon, A. Elisseeff. An Introduction to Variable and Feature Selection. Journal of Machine Learning Research, 3:1157-1182, 2003.