A Game Theoretic Analysis of the Twitter Follow-Unfollow Mechanism

Jundong Chen[1(✉)], Md Shafaeat Hossain[2], Matthias R. Brust[3], and Naomi Johnson[1]

[1] Department of Math and Computer Science, Dickinson State University, Dickinson, ND 58601, USA
{jundong.chen,naomi.johnson}@dickinsonstate.edu
[2] Computer Science Department, Southern Connecticut State University, New Haven, CT 06515, USA
hossainm3@southernct.edu
[3] Interdisciplinary Centre for Security, Reliability and Trust (SnT), University of Luxembourg, Esch-sur-Alzette, Luxembourg
matthias.brust@uni.lu

Abstract. Twitter users often crave more followers to increase their social popularity. While a variety of factors have been shown to attract the followers, very little work has been done to analyze the mechanism how Twitter users follow or unfollow each other. In this paper, we apply game theory to modeling the follow-unfollow mechanism on Twitter. We first present a two-player game which is based on the Prisoner's Dilemma, and subsequently evaluate the payoffs when the two players adopt different strategies. To allow two players to play multiple rounds of the game, we propose a multi-stage game model. We design a Twitter bot analyzer which follows or unfollows other Twitter users by adopting the strategies from the multi-stage game. We develop an algorithm which enables the Twitter bot analyzer to automatically collect and analyze the data. The results from analyzing the data collected in our experiment show that the follow-back ratios for both of the Twitter bots are very low, which are 0.76% and 0.86%. This means that most of the Twitter users do not cooperate and only want to be followed instead of following others. Our results also exhibit the effect of different strategies on the follow-back followers and on the non-following followers as well.

Keywords: Social network · Game theory · Machine learning Twitter classification · Twitter bot

1 Introduction

On the Twitter platform, a user can follow and can be followed by other users. In the early stage, Twitter allowed users to follow as many accounts as possible. Many Twitter users abused this and hoped to increase the number of followers through following thousands of users instead of creating engaging content.

L. Bushnell et al. (Eds.): GameSec 2018, LNCS 11199, pp. 265–276, 2018.
https://doi.org/10.1007/978-3-030-01554-1_15

Therefore, Twitter set up a limit for the number of accounts users could follow. The number of accounts that a Twitter user can follow cannot be 10% more than the number of followers, and also must be less than 2,000. In 2015, Twitter changed this limit to 5,000.

A follower's count is one of the three measures which indicate a Twitter users' popularity and prestige [1]. Researchers have been investigating the variables which effect the follower behavior of online social networks (OSN). Hutto et al. found that social behavior, message content, and network structure have different effects on determining other Twitter users to follow a Twitter user [2]. Liu et al. built a model for inferring the different speed of follower growth of different types of users on a microblog platform (Weibo) [3]. Mueller et al. integrated multiple predictors from the profile information of a Twitter user to predict the increase of the follower count [4].

Some researchers use Twitter bots to manipulate Twitter accounts to attract followers to create influential Twitter accounts. Those Twitter bots implement the functions of a regular Twitter account which is managed by a real user. Such functions include follow, unfollow, and post tweets, etc. A Twitter bot is a type of automated program which controls a Twitter account via Twitter API [5]. Messias et al. found that a Twitter account operated by a Twitter bot is capable of becoming influential by mimicking a real Twitter user through simple strategies, such as following back the followers and posting tweets about trending topics [6].

Game theory has been applied to model the influence from the interactions between OSN users on the privacy settings. Chen et al. modeled privacy settings of online social networks by a two-player game and an evolutionary game, and investigated the effect of network connectivity and attribute importance on the users' profile disclosure [7,8].

For this paper, we developed two game theoretic models to analyze the Twitter follow-unfollow mechanism. One is a two-player game, which is called Twitter follower's dilemma. The other one is called multi-stage follow-unfollow game, which allows players to play the game multiple rounds. Then, we designed two Twitter bot analyzers[1] which can adopt the strategies derived from the game models. Subsequently, the Twitter bot analyzers collect the response from other Twitter users when different strategies are adopted. Our approach not only explores the dynamics of the users when we follow them, but also discovers the impact of the adopted strategies on the non-following users. We call the users that follow us back the follow-back followers. The non-following followers mean the users we do not follow but they still follow us.

The remainder of this paper is as follows. In Sect. 2, we derive the two-player follow-unfollow game from the Prisoner's Dilemma game, and subsequently the multi-stage follow-unfollow game. In Sect. 3, we explain the method for classifying the collected Twitter users. The process of data collection, the experiment

[1] The two Twitter bot analyzers follow the same steps, except that Twitter bot 1 takes one more step, which is favoriting the tweets posted by other Twitter users. This is to investigate the effect of favoriting tweets on the number of followers.

design, and the algorithm that the Twitter bot utilizes in the multi-stage game are elaborated in Sect. 4. We present the results and the discussion in Sect. 5. We conclude this paper in Sect. 6.

2 Our Models

2.1 Twitter Follower's Dilemma

Our approach to model the Twitter followers' dynamics is inspired by the Prisoner's Dilemma [9]. The Prisoner's Dilemma models a situation with two completely rational individuals who might not cooperate, even if it is in their best interests to do so. It provides a framework for us to understand a balance lingering between cooperation and competition. In our game, there are two players which are the two Twitter users, user A and user B. In each step each user can choose between two strategies, "follow" and "unfollow". The goal of each player is to achieve high social popularity [10], which means to have as many followers as possible.

The payoff matrix for the Twitter follower's dilemma game is shown in Table 1. There are 4 cells in the matrix. Each cell has a tuple which represents the payoff for user A and user B, respectively. Therefore, we have 4 different combinations according to different strategies adopted by the two users, which are **(follow, follow)**, **(follow, unfollow)**, **(unfollow, follow)**, and **(unfollow, unfollow)**. We can summarize all these combinations into 3 different cases, because (follow, unfollow), (unfollow, follow) are symmetric.

Case I: This case refers to **(follow, follow)**. After one user follows the other one, and the other one also responds with a "follow" strategy, then each one receives a modest payoff, which is denoted by 2. This is because each user is followed by the other one but still needs to invest one count of "following".

Case II: This case refers to **(follow, unfollow)** or **(unfollow, follow)**. When one user follows the other one but the other one has not responded with the "follow" strategy, then the user being followed gets more benefit because this user can follow more accounts because of getting this following. In this case, we say that a user with an "unfollow" strategy achieves the highest payoff denoted by 3. However, the other user has the lowest payoff denoted 0. This is because one user invests one count of following but this following ends up with no increase in the number of followers, and this investment is in vain.

Case III: This case refers to **(unfollow, unfollow)**. This case may happen before or after these two users interact. Before they interact, no one takes any action, which means "unfollow" for each one. After one user follows the other one and later finds that the other one has no response, then this user decides to disconnect with the other one. In this case, each user receives a payoff of 1, which means no one reaches the highest payoff.

In this game we assume that one user decides to adopt any strategy by only considering the payoff from the social popularity. We know that in some situations we can already benefit from only following an account. For example, if one user is a fan of a celebrity from following the celebrity's account, this user receives the status update or some interesting activities. Or, if we follow some Twitter account of a news website, we receive interesting news or stories.

Table 1. Payoff matrix for the follower's Dilemma game.

		Twitter User B	
		Follow	*Unfollow*
Twitter User A	*Follow*	$(2, 2)$	$(0, 3)$
	Unfollow	$(3, 0)$	$(1, 1)$

2.2 Revised Twitter Follower's Dilemma

After considering the follower's benefit of receiving news, we can revise the game in Sect. 2.1, we obtain the following payoff matrix as shown in Table 2. We use N to represent the benefit from receiving news. In this payoff matrix, since we are using a Twitter bot as the player and the Twitter bot will not read the news received from other Twitter users, news is not considered as a benefit for the Twitter bot player.

Table 2. Payoff matrix for the revised Twitter follower's Dilemma.

		Twitter User	
		Follow	*Unfollow*
Twitter Bot	*Follow*	$(2, 2 + N)$	$(0, 3)$
	Unfollow	$(3, 0 + N)$	$(1, 1)$

2.3 Multi-stage Follow-Unfollow Game

One Twitter bot in our experiment plays multiple rounds of games with other Twitter accounts by taking follow or unfollow strategies in turns. This process is modeled as a multi-stage game as shown in Fig. 1.

In Fig. 1, $P1$ represents player 1 which is our Twitter bot, and $P2$ represents a group of other Twitter users which is player 2 in this multi-stage game. In this game, $P1$ at first follows all the Twitter users. Some of the Twitter users follow, and others unfollow. For those Twitter users who do not follow, after waiting for a period of time our Twitter bot gives up on them and unfollows them. For those Twitter users who follow our Twitter bot, we play more rounds of the game. After they follow us, our Twitter bot unfollows them with the intent of maximizing the payoff. Some Twitter users may notice that they are unfollowed and as a countermeasure they unfollow our Twitter bot. Other Twitter users may still follow. For those Twitter users who adopt the strategy of "unfollow" as the countermeasure, our Twitter bot attempts to regain them and follow them again. Some users may follow back again, however, other users may already lose their trust to our Twitter bot and never follow back.

The expected payoff of the Twitter bot is calculated by

$$U = 3\alpha\beta + 2 \cdot \alpha(1-\beta)\gamma + 1 \cdot \alpha(1-\beta)(1-\gamma) + 1 \cdot (1-\alpha) \tag{1}$$

where α, β, and γ represent the ratio of users who adopt a follow strategy at different stages, which are denoted in the parenthesis behind each strategy.

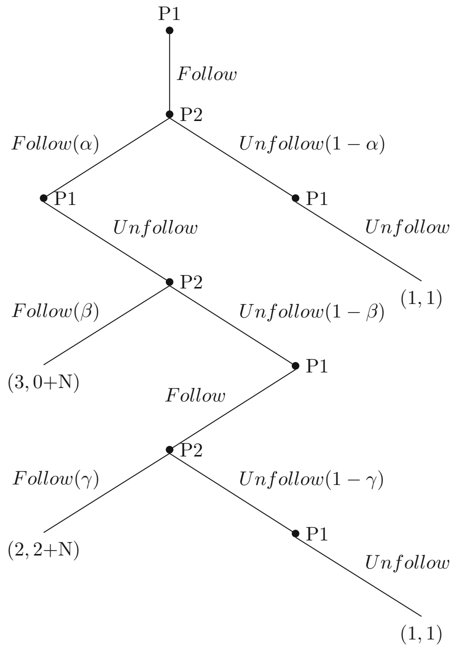

Fig. 1. Extensive form of the multi-stage follow-unfollow game.

3 Twitter User Classification

We use a machine learning method presented by Deshpande on PyCon France 2016 [11] to classify the Twitter users based on the tweets posted by each user.

In our experiment, we choose 8 typical categories, which include "Tech", "Business & CEOs", "Entertainment", "Science", "Fashion, Travel & Lifestyle", "Sports", "Music", and "Politics" as shown in Table 3.

Table 3. Category IDs and the corresponding category names.

Category ID	Category name
Category 0	Tech
Category 1	Business & CEOs
Category 2	Entertainment
Category 3	Science
Category 4	Fashion, travel & lifestyle
Category 5	Sports
Category 6	Music
Category 7	Politics

4 Experiment

In our experiment, we use a Twitter crawler to collect Twitter users' ids and retrieve tweets for all of these users. A Twitter classifier assigns all the users into different categories. Then, our Twitter bot plays game with the Twitter users in different categories. We record the list of friends[2] and that of followers for the Twitter bot account over time. The structure of the experiment is given in Fig. 2.

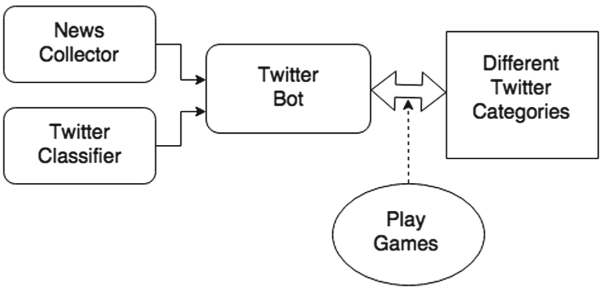

Fig. 2. Design of the Twitter bot analyzer.

4.1 Procedure of the Experiment

We proceed with the experiment by the following steps.

Step 1: Construct User Lists.
We use a Twitter crawler supported by Twitter API [12] to collect the user dataset. In this dataset, we apply the snow ball sampling technique [13] to collect

[2] The two terms, friends and followees, are interchangeable on Twitter. If we follow one user, we can call that user as a friend or followee of our Twitter account.

Twitter users' ids. Each time running the crawler, we start from a different Twitter account, which is called "seed". We collect the user lists by selecting different seeds at different locations in the world. Here, the different locations correspond to different geographic coordinates.

In total, we have collected 11,349 Twitter ids, and about 1,000 tweets for each Twitter user on average. We separate the Twitter ids in each category into two groups for two Twitter bots. Then, we mix the Twitter ids from different categories into one file and shuffle them. This is to ensure that each user is randomly assigned to each Twitter bot and also guarantee that the users in each category are equally assigned to the two Twitter bots.

Step 2: Post News from Different Sections.

Twitter bots follow the users in different lists and then post tweets with the news from different news sections from ABC news.

In order to make tweet contents attractive to different types of people, we post different types of tweets. We classify the Twitter accounts into 8 different categories as in [11,14], which are listed in Table 3. Everyday we crawl news on the website of ABC news[3]. There are only 5 sections of news which match the interests of 5 different types of twitter users, which are "Technology", "Entertainment", "Lifestyle", "Sports", and "Politics". We post that message from that sharing link obtained by clicking the Tweet share button.

Step 3: Twitter Bots Play a Game with Twitter Users.

As shown in Table 1, Twitter bots have two strategies to adopt. Depending on the different strategies taken by different users, the Twitter bots respond with different strategies.

The Twitter bot has to follow other Twitter accounts first to attract them in order to increase the number of followers as a consequence. After being followed, this bot will unfollow that follower to spare the quota of followings and spend this number to follow another new account. After this bot unfollows an account, that Twitter account may take a countermeasure to unfollow the bot. Then, this bot follows back again. The follow and unfollow strategies may be adopted by the bot and a Twitter account by turns in a couple rounds. We use a multi-stage game to model this process as shown in Fig. 1.

Step 4: Collect the Data About the Dynamics of Strategies of the Users.

Every day we check the followers and followees of our Twitter accounts. Then we draw a trend curve for each of our Twitter accounts to show the changes of the number of followers over time.

4.2 Twitter Bot Analyzer in the Multi-stage Game

We present the pseudocode in Algorithm 1 which describes the workflow of a Twitter bot in the multi-stage game. The whole process is divided into two

[3] http://abcnews.go.com.

phases. In all the phases, our Twitter bot keeps posting tweets of the news from different news sections collected from ABC news. After the Twitter bot starts following other users, we save the friends and followers of the Twitter bot account into separate files each day. The purpose of the first phase is to attract attentions of other Twitter users. In the second phase, the Twitter bot plays the game with other Twitter users by taking the strategies described in the model as shown in Fig. 1.

The first phase in the whole process is to follow Twitter users, and like the tweets from those users. With the limits from Twitter, we only follow 1000 accounts in one day, and favorite 1200 tweets per day and one tweet per minute. Twitter prohibits any aggressive following behavior, therefore we follow Twitter users with the amount below the limit.

In the second phase, we keep tracking the followers and unfollower in different stages and assign them into different sets. After waiting a period of time that our Twitter bot follows all the Twitter users that we have collected, some users follow back, and others do not. The followers are assigned into set S, and unfollowers into set S'. The Twitter bot unfollows all of them, which is a strategy decided in the algorithm. After passing through a date range from d_3 to d_4, the users that are still followers are assigned into set S_1, and unfollowers into S_2. The bot follows the users in S_2 trying to regain their trust. After waiting a period, some users in set S_2 follow back and others do not. Then, we save the followers from the set S_2 to S_{21}, and unfollowers to S_{22}.

The set notations in the model and the algorithm are depicted in Fig. 3.

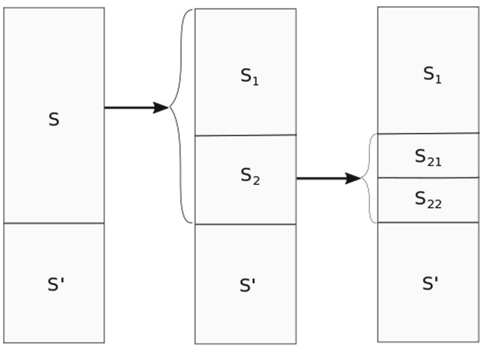

Fig. 3. Set structure diagram at different stages of the game.

5 Results and Discussion

The Twitter bot analyzers keep tracking the followers of our Twitter bots once we start the experiment as described in Algorithm 1. Figure 4a and b show the number of followers of the two Twitter bots changes over time. The figures only show the records after the Twitter bots finish following all the assigned Twitter

Algorithm 1. Twitter bot analyzer in the multi-stage follow-unfollow game

Input: Twitter accounts in 5 different categories
Output: Record of friends and followers of the twitter bot in each day

$countDays = d_0 - 1$
for each day in a date range $[d_0, d_1]$ **do**
    countDays++
    Tweet news from different news sections
    if Total number of followings less than 5000 **then**
        Follow each of the Twitter users, 1000 users per day
    end if
    Retrieve friends and followers ids
end for
for each day in a date range $(d_1, d_2]$ **do**
    countDays++
    Tweet news from different news sections
    for every minute in a total of 20 hours **do**
        Favorite a tweet for each of the Twitter users
    end for
    Retrieve friends and followers ids
end for
for each day in a date range $(d_2, d_3]$ **do**
    countDays++
    Retrieve friends and followers ids
end for
Save all followers in set S
Save all unfollowers in set S'
Unfollow the users in set S'
Unfollow the users in set S
for each day in a date range $(d_3, d_4]$ **do**
    countDays++
    Retrieve friends and followers ids
end for
Save the followers from set S to S_1
Save the unfollowers from set S to S_2
Follow the users in S_2
Retrieve friends and followers ids
Save the followers from set S_2 to S_{21}
Save the unfollowers from set S_2 to S_{22}
Unfollow the users in set S_{22}

users. Because of the limit from Twitter, each Twitter bot can only follow up to 5,000 Twitter accounts in total and about 1,000 per day. It takes 5 days to follow about 5,000 Twitter users. Day 0 in Fig. 4 means the $5th$ day after the Twitter bots start following Twitter users.

Originally, there are 5,000 Twitter ids in the list for each of the Twitter bots. However, in fact, Twitter bot 1 follows 4,981 users, and Twitter bot 2 follows

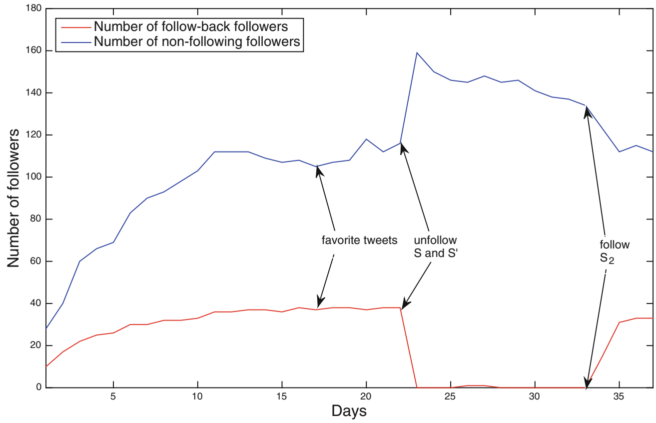

(a)

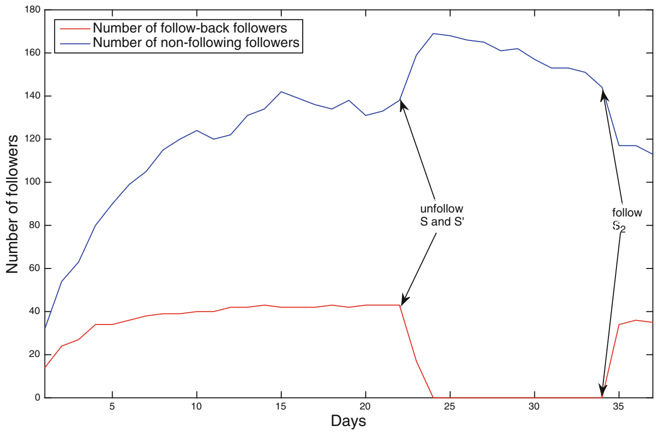

(b)

Fig. 4. The number of followers for each Twitter bot changes over time. (a) Twitter bot 1, (b) Twitter bot 2.

4,980 users. This is because some of the accounts in the list are suspended or not used after we build the list, and no one can follow them.

To better show the exact values, we use Table 4 to list the number of users encompassed in different sets. The definitions of the sets are given in Sect. 4.2 and depicted in Fig. 3. The size of S and that of S' together are equal to the total number of Twitter users that a Twitter follows at the beginning of the experiment. Set S is divided into S_1 and S_2 depending on if they unfollow after the Twitter bots unfollow them. If they unfollow, then they are assigned to set S_2, otherwise remain in set S_1.

The first observation is that Figs. 4a and b exhibit almost the same change pattern for the number of followers. In each figure, there are two curves. One curve is for the number of followers who follow back, which is denoted as "number of follow-back followers". The other one is for the number of followers that our Twitter bots never follow, with the legend of "number of non-following followers". We find that at first both of the curves increase as time elapses. The number of non-following followers increases until the Twitter bots unfollow the Twitter users in set S and S'. In all the follow-back followers which are in set S, if the Twitter bots unfollow them, they unfollow our Twitter bots immediately. This is why in the table the size of the set S_1 for each bot is zero. This means that for these users they adopt the strategy "follow" if the adversary has a strategy "follow", and respond with an "unfollow" strategy to an "unfollow" strategy.

An interesting observation is that although Twitter bot 1 favorites the tweets of other users, the number of followers still shows almost the same change pattern. For both of the bots, the number of the non-following followers keeps increasing for a short period and then drops. This means that favoriting the tweets of other users has little effect on the change pattern of the number of followers.

The second observation is from Table 4, Twitter bot 1 only gains 38 follow-back followers after following 4,980 users, and bot 2 gets 43 follow-back followers after following 4,981 users. The follow-back ratio is 0.76% and 0.86% for Twitter bot 1 and 2, respectively. Both ratios are very low. This is not coincident and is explained by our Twitter follower's dilemma game model. Most of the Twitter users do not cooperate and they only want to be followed instead of following other users.

Table 4. The sizes of different sets at different phases of the game.

Set		S	S'	S_1	S_2	S_{21}	S_{22}
Size	Twitter bot 1	38	4,942	0	38	33	5
	Twitter bot 2	43	4,938	0	43	35	8

6 Conclusions

In this paper, we analyze the mechanism on Twitter about users following or unfollowing others. We propose a two-player game, which is called a Twitter follower's dilemma. In this game, each player has two strategies: follow and unfollow. Then, we design a multi-stage follow-unfollow game.

We also create two Twitter bot analyzers. The two analyzers prove that the finding from one analyzer is not coincident and furthermore, investigate the effect from favoriting tweets on the number of followers.

Two Twitter bots show the same change pattern for the number of followers. Another finding is that for all the follow-back followers, if we unfollow them, they unfollow us as a countermeasure. Our results show that the follow-back ratios are very low. Our results also show that favoriting tweets of other users has little effect on the number of followers. As a by-product, our results exhibit the change pattern for the number of non-following followers.

The approach presented in this paper provides a way to analyze and investigate the Twitter follow-unfollow mechanism and helps to optimize the design of a social network platform.

References

1. Cha, M., Haddadi, H., Benevenuto, F., Gummadi, K.P.: Measuring user influence in Twitter: the million follower fallacy, In: ICWSM (2010)
2. Hutto, C., Yardi, S., Gilbert, E.: A longitudinal study of follow predictors on Twitter. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI 2013, pp. 821–830. ACM, New York (2013)
3. Liu, L., Song, D., Tang, J., Liao, L., Li, X., Du, J.: FriendBurst: ranking people who get friends fast in a short time. Neurocomputing **210**(C), 116–129 (2016)
4. Mueller, J., Stumme, G.: Predicting rising follower counts on Twitter using profile information. In: Proceedings of the 2017 ACM on Web Science Conference, WebSci 2017, pp. 121–130. ACM (2017)
5. Jajodia, S., Wang, H., Gianvecchio, S., Chu, Z.: Detecting automation of Twitter accounts: are you a human, bot, or cyborg? IEEE Trans. Dependable Secur. Comput. **9**, 811–824 (2012)
6. Messias, J., Schmidt, L., Oliveira, R., Benevenuto, F.: You followed my bot! Transforming robots into influential users in Twitter. First Monday **18** (2013)
7. Chen, J., Kiremire, A.R., Brust, M.R., Phoha, V.V.: Modeling online social network users' profile attribute disclosure behavior from a game theoretic perspective. Comput. Commun. **49**, 18–32 (2014)
8. Chen, J., Kiremire, A.R., Brust, M.R., Phoha, V.V.: A game theoretic approach for modeling privacy settings of an online social network. EAI Endorsed Trans. Collab. Comput. **1**(1–5), e4 (2014)
9. Surhone, L., Timpledon, M., Marseken, S.: Prisoner's Dilemma: Game Theory, Merrill M. Albert W. Tucker, Framing Device, Experimental Economics. Betascript Publishing (2010)
10. Moses, J.: The follower's dilemma, how game theory explains Twitter's most important feature. https://www.linkedin.com/pulse/followers-dilemma-how-game-theory-explains-twitters-most-moses. Accessed 1 Feb 2016
11. Deshpande, D.: Twitter user classification with gensim and scikit-learn. In: The Python Conference, France (2016)
12. Zafarani, R., Abbasi, M.A., Liu, H.: Social Media Mining: An Introduction. Cambridge University Press, Cambridge (2014)
13. Kay, M.: Generating a network graph of Twitter followers using python and networkx. http://mark-kay.net/2014/08/15/network-graph-of-twitter-followers/. Accessed 15 Aug 2014
14. Chen, J., Li, H., Wu, Z., Hossain, M.S.: Sentiment analysis of the correlation between regular tweets and retweets. In: Proceedings of the 16th IEEE International Symposium on Network Computing and Applications, pp. 1–5 (2017)