Characterizing the Underlying Social Network of Prescription Drug Abusers

M.Tech Project Report Submitted in Partial Fulfillment of the Requirements for the Degree of

Master of Technology
in
Computer Science & Engineering

by

Sequeira Ryan Thomas
(1611CS13)

under the guidance of

Dr. Joydeep Chandra

DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING

INDIAN INSTITUTE OF TECHNOLOGY PATNA
PATNA - 801106, BIHAR

CERTIFICATE

*This is to certify that the work contained in this thesis entitled "**Characterizing the Underlying Social Network of Prescription Drug Abusers**" is a bonafide work of **Sequeira Ryan Thomas (Roll No. 1611CS13)**, carried out in the Department of Computer Science and Engineering, Indian Institute of Technology Patna under my supervision and that it has not been submitted elsewhere for a degree.*

Supervisor: **Dr. Joydeep Chandra**

Assistant Professor,

Department of Computer Science & Engineering,

Indian Institute of Technology Patna, Bihar.

July, 2018

Patna.

Abstract

Prescription drug abuse is fast becoming a menace as youths across different societies are being severely affected. The Twitter platform is being actively used for discussions where drug abusers are sharing their experiences and glorifying drug abuse. Such discussions not only help the drug abusers in rationalizing their habits but also act as social advertisements that further the spreading of such abuses. Hence analyzing the user engagement in drug abuse tweets can be key to understand the role of social media in spreading of the menace. We perform a large-scale study of the Twitter follower network involving around 0.28 *million drug abusers to characterize the user engagement and spreading of drug abuse tweets across the network. Our observations reveal the existence of a very large giant component involving* 93% *of the users (drug abusers) that facilitates the spreading of such messages. Further observations indicate the presence of few large cascades of user engagement, with multiple users playing key roles in the spreading. Moreover, our observations also reveal a collective phenomenon, involving a large set of active fringe nodes (with a small number of follower and following) along with a small set of well-connected non-fringe nodes that work together towards such spread. We also observe the engagement of users with respect to drugs like Vicodin, Percocet, OxyContin, Lortab and Dilaudid. The engagement probability continues to remain high with increasing exposure to such tweets, thereby indicating that vulnerable candidates slowly get engaged through discussions in social media. Finally we attempt to uncover the promoters in this network. Our approach identifies these promoters with nearly* 50% *accuracy that is significantly higher as compared to other baselines. We also show that* *identifying the promoters using our approach can also help in the early identification of potential drug abusers. Using experiments on a temporal data we show that our approach is able to predict the potential abusers with more than 70% accuracy.*

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Recent reports by *The Economist*[1] highlight the spread of prescription drug abuse reaching almost epidemic proportions in the USA. Data presented by *CDC*[2] reveal that majority of the drug abusers are teenagers between $18-22$ years of age. Further, there is an estimated count of $190,000$ premature drug-related deaths and the majority is due to the use of opioids[3]. The popularity of social media like Twitter is attracting illegal drug agents to leverage the same for promoting drugs by reaching out to vulnerable users. Keeping in view the spread and impact of drug abuse we, therefore, need to have a deeper understanding of how social media is playing a role in promoting drug menace.

We focus our attention on the Twitter platform, that is one of the key media used to spread information related to drugs. Several background works exist that have highlighted the role of Twitter in the sale of illicit drugs [30] and promotion of drug abuse [31]. Possible surveillance strategies for identifying such retailers and drug abusers have also been well explored [25]. However, an important aspect that needs to be carefully investigated is the networked affect of Twitter that may amplify the spread of drug abuse among users. Preliminary studies of Twitter users in [20] reveal the existence of social circles (densely

[1] https://www.economist.com/blogs/graphicdetail/2017/03/daily-chart-3
[2] http://wonder.cdc.gov/mcd.html
[3] https://www.cdc.gov/nchs/data/health_policy/monthly-drug-overdose-death-estimates.pdf

connected neighbor set of a user) around certain active users who tweet frequently mentioning different effects specific drugs produce. Further, their work also suggests that such circles fulfill the need of the drug abusers to connect among themselves and observe the sentiments of other users on the subject of interest. Although these observations are insufficient to establish whether such active circles can influence non drug-abusers towards abuse, however, from the Health Communication Media Choice (HCMC) model, proposed in [16], it can be inferred that non-active abusers (who abuse drugs but do not tweet) may easily rationalize their drug abuse behavior from such discussions that primarily glorify drug abuse. Consequently, cascades of user engagement, if formed through such discussions, would volume up as social advertisements that would not only influence vulnerable users towards drug abuse but would also complicate rehabilitation processes.

We initially propose a technique that uncovers around 0.28 million unique users who were engaged in either self-reporting or promoting drug abuse through tweets. The enormity of this number reflects the huge role being played by the social media in promoting prescription drug abuse on Twitter. The *follower-followee* relation among these users is used to create a network with directed edges that we term as *Prescription Drug Abuse Network* (PDAN). Further investigation reveals the existence of a large connected component in the PDAN involving 93.24% of the nodes. A bow-tie representation of this giant component reveals a heavy core structure with 86.77% of the nodes in the largest strongly connected component, and only 10.2% and 1.5% nodes in the IN and OUT components respectively. We also observe that PDAN has high clustering coefficient and reciprocity between links. All these network properties indicate the network structure is amenable to the large-scale spread of information. The spread, however, depends on the activeness of the users (frequency of engagement) as well as their position in the network. We assess the activeness of the users and found that around 3-4% of the users are active and 10% among them have central positions in the network. However, the reach of these active nodes is considerable, they cover a substantial part of the network.

The detailed study of the nature of the network, as well as activeness of the users, facilitate us to understand the formation of cascades over such network. We discover around $33,628$ cascades, several of them with sizes reaching to thousands and extending over several hops. The network study of such cascades reveal high structural virality, that is the cascades are not driven by a single *important* node which in turn makes them difficult to control. Guided by the metrics associated with social contagion processes [35] (discussed in detail later), we observe that for certain abused drugs like Vicodin and Percocet, that are opioid painkillers and found significantly high mentions in the tweets, the users are highly prone to exposure and a large number of users can get influenced at the same time. The study of cascades with respect to the position of users reveal that around 40% of all cascades are initiated by fringe nodes (nodes with very low number of followers as well as followings). Observations indicate that the cascades initiated by these fringe nodes propagate one or two hops through other fringe nodes before expanding through non-fringe nodes. The phenomenon indicates that organic collaboration among a large set of nodes is largely responsible for the emergence of cascade thus it may be difficult to control the cascades by eliminating a few targeted nodes.

We make an effort to solve this problem by identifying promoters in this network. We use a combination of content as well as network analysis based approach in identifying the promoters influencing drug abuse activities either directly or through implicit means. A directed edge from a node indicates the possible influence of that node on the node at the other end of the edge. The weight value provides a measure of this personal influence calculated based on their content similarity and the temporal characteristics of the tweets sent by the corresponding users. The promoters are identified and ranked based on a node impact score derived using a modified Katz centrality [26] based approach. The node impact scores are calculated based on the weight value of the links and are calculated using a lightweight random-walker based algorithm. We validate the proposed approach using a validation data set and compare the precision, recall and accuracy of our approach with several baselines. Experimental results indicate that the proposed approach outperforms other baseline approaches by a significant margin. Subsequently, we show that the proposed influence measure used for identifying the masquerades can be used for early identification of the vulnerable and potential drug abusers. Experimental results on a temporal dataset indicates that the proposed approach can identify potential abusers with nearly 70% accuracy.

Chapter 2

Related work

A plethora of recent works use social media to gather information and in turn, provide solutions to various issues related to health. For example, social media has played an important role in providing rich information for inferring mental health conditions (especially depression [11], mood instabilities [14] and suicidal risks [12]), as well AS lifestyle-related conditions like overeating, alcoholism and smoking [44]. Technological approaches are being leveraged for addressing key issues like the early prediction of such diseases [11], increased support and service engagement [36] and decrease the duration of untreated disorders [10]. These works provide a direction to the key issues, with respect to the psychological problems (that includes drug abuse problem), that require immediate attention.

Recently, prescription drug abuse is receiving increasing attention due to its significant spread and the casualties involved [24, 27]. Social media has proved to be an important resource in obtaining abuse-related information, especially contents that reveal drug abuse behavior of users [40, 33]. Authors in [40] proposed a technique to automatically identify such abuse tweets based on supervised classification and natural language processing. Similarly, in [7], the authors used a large set of keywords for searching the contents and used content similarity-based approaches for identifying abuse tweets. Strategies and guidelines for creating annotated datasets required for classification have been discussed in [39].

In [27, 1], the authors prepared an enriched corpus of pharmacovigilance curated from Twitter messages. Several unsupervised techniques have also been proposed to identify drug abuse tweets from tweet stream [24]. Authors in [15] applied topic modeling approach to automatically disambiguate hashtags based on their topical context to classify abuse tweets. These works help in identifying the drug abuse tweets from the tweet stream and thus provide an opportunity to explore the other key facets to the drug abuse problem, including ways to mitigate their spread in the society. This research attempts to work toward this direction; we undertake a detailed investigation to understand the spread of drug abuse messages through Twitter and possibly identify promoters of prescription drugs.

Very few works exist that discuss the microscopic behavior of the abusers and their role in spreading of drug abuse tweets in the network. In [20], the authors pointed out the influence of neighbors (network effect) on the participation of users in discussions related to drug abuse. They observed the presence of social circles in which active users (users who frequently discuss about abuse of specific drugs in Twitter) are likely to be surrounded by users who also participate in similar discussions, exhibiting high content correlation among them. While this work provides preliminary insights about the tweeting behavior among the drug abusers, revealing the existence of a possible group phenomenon; to the best of our knowledge, no other work has attempted to take a closer look into the spread of drug abuse discussions on social media platforms. However, the study of the propagation of different tweet contents to understand human behavior has been a focus area across various topical domains. One of the earlier works on information flow in Twitter [47] showed the presence of a few elite users in Twitter who generate a majority of the contents that are consumed by ordinary users. Several other works have also investigated the user characteristics and role of influential users in information propagation across social networks [3, 4]. Subsequent empirical works have examined several other factors like the underlying network of the users [46, 5, 9], the user characteristics [34, 38, 18] and the role of content [43, 45] in such propagation. Certain works have investigated the effects of both users and contents in information propagation [22, 37]. Since the dynamics of information propagation across networks vary with topics and contents, this motivates the need to investigate the spreading behavior of drug abuse tweets by the users and explore the role played by the network and user characteristics in such a spread.

Spreading of various social behavior has been investigated in several works, like cessation of smoking [13], online sharing [41], and political controversies [35]. These works highlight the importance of collective dynamics, often modeled through a complex contagion phenomenon, in the spreading of social behavior. As these works can eventually help in controlling viral spread when such propagation is not desired (like in case of drug abuse tweets), there is a need to look into the generation of drug abuse tweets through the prism of such models. As there is still a wide gap in understanding the characteristics of the social network of drug abusers along with their network roles that influence *spreading* of drug abuse contents, we believe this research would contribute in filling this gap.

In the next chapter, we describe the dataset used for this study and network preparation.

Chapter 3

Constructing the Prescription Drug Abuse Network

3.1 Dataset

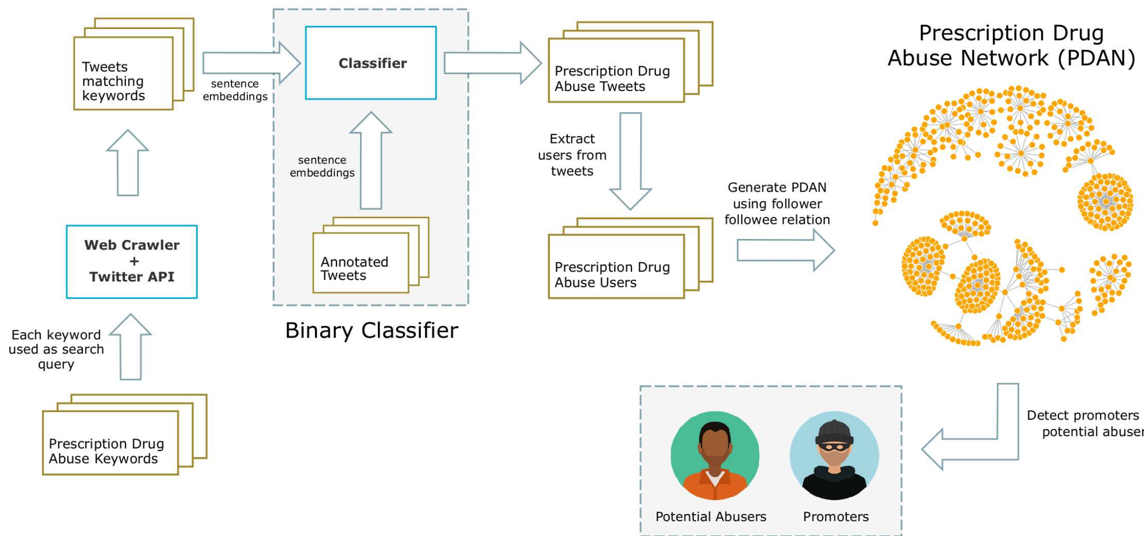

Fig. 3.1: Steps of PDAN formation.

In this section, we provide the detailed description of the tweet dataset along with the data collection methodology and the preprocessing techniques used. We subsequently

Generic names	Brand names
oxycodone	OxyContin, Percodan, Percocet
hydrocodone	Vicodin, Lortab, Lorcet
diphenoxylate	Lomotil
morphine	Kadian, Avinza, MS Contin
codeine	-

Generic names	Brand names
fentanyl	Duragesic
propoxyphene	Darvon
hydromorphone	Dilaudid
meperidine	Demerol
methadone	-

Table 3.1: List of generic and brand names of prescription opioids medically used to treat pain.

describe the classification technique used to identify tweets that are promoting or reporting prescription drug abuse (the corresponding promoters are henceforth termed as drug abusers). Based on the follower-followee relation among these drug abusers, a network termed as *Prescription Drug Abuse Network* is created. The steps followed to form the prescription drug abuse network is pictorially represented in figure 3.1.

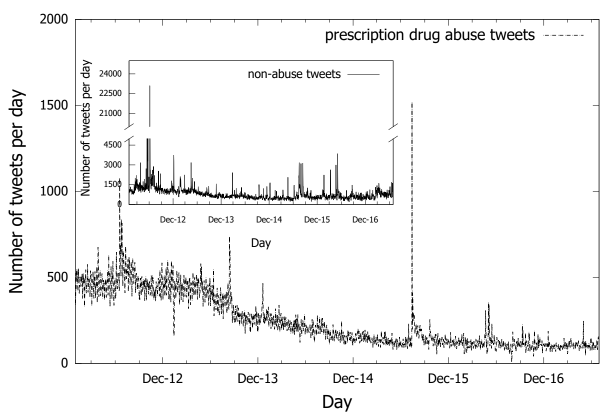

Fig. 3.2: Timeseries of tweet activity based on the tweets collected using keyword search from 2012 to 2017 using web-based crawler[21]. The figure shows the number of tweets per day classified as prescription drug abuse and the inset figure shows the number of tweets per day classified as non-abusive tweets. The data is plotted as time series data to smoothen the local irregularities and highlight trends.

The number of tweets in the tweet stream that are self-reporting drug abuse is very sparse. In contrast to 500 *million tweets posted on Twitter each day*[1] we observed that, on an average only 1,070 tweets per day contained certain prescription drug abuse keywords,

[1]https://www.omnicoreagency.com/twitter-statistics/

and further an average of only 246 tweets per day (i.e. 20% of collected tweets containing drug abuse keywords) were classified as prescription drug abuse tweets (figure 3.2). The rest of the tweets were related to spreading awareness about the drug abuse problem, promote rehabilitation to drug abusers or contained these keywords in a context not related to drug abuse. Hence to overcome this data sparsity, we use a technique that is outlined next.

3.1.1 Data collection

The data collection steps can be briefly described as follows:

1. We prepare a set of drugs names (a non-exhaustive list is shown in table 3.1) that have been marked and listed for abusive use in the past by National Institute on Drug Abuse (NIDA)[2]. The generic and brand names of these drugs were used as search *keywords* for collecting the drug-related tweets using the web-based crawler, "*Get Old Tweets*" [21]. This gave us all the searchable tweets from January 2012 to July 2017 containing the drug names. Retweets are not retrieved using this web-based crawling approach and are obtained using a different Twitter API[3].

2. As only a limited information about the tweets were being provided by the web-based crawler, we used the "*tweet id*" of the returned tweets to further query and extract the complete tweet information using the Twitter API[4].

Using this approach, we collected more than 2 million drug-related tweets. However, as stated earlier, we observed that this collection of tweet set includes both kinds of tweets: those promoting drugs or reporting drug-abuse as well as those spreading awareness or rehabilitation and treatment information, that we term as non-abuse tweets. Hence, we apply an SVM based machine learning technique to identify the drug-abuse tweets which is detailed in the next section.

[2]https://teens.drugabuse.gov/drug-facts/prescription-pain-medications-opioids
[3]https://api.twitter.com/1.1/statuses/retweets/:id.json
[4]https://api.twitter.com/1.1/statuses/show.json

3.1.2 Classification of Drug Abuse Tweets

Classifier	N-grams as features			sentence embeddings and handcrafted features		
	Abuse F_1	Non-Abuse F_1	Accuracy	Abuse F_1	Non-Abuse F_1	Accuracy
Naive Bayes	0.752	0.701	72.87	0.758	0.735	74.718
SVM	0.787	0.759	77.42	**0.857**	**0.846**	**85.16**
Random Forest	0.842	0.794	82.12	0.814	0.806	81.03
Logistic Regression	0.701	0.694	69.75	0.811	0.803	80.70
Ensemble	0.740	0.613	68.90	0.782	0.770	77.60

Table 3.2: F_1 score measured for the two classes (Abuse and Non-Abuse) using 10-fold cross-validation. sent2vec[32] with handcrafted features significantly outperforms n-gram classifiers without any feature engineering and relatively smaller feature vectors.

One of the key requirements in characterizing the drug abusers is to classify the tweets based on whether they promote/report prescription drug abuse or not. On manual categorization of a sample of tweets (containing drug names), we observed, apart from tweets that promote sales and report drug abuse, a significant proportion of the tweets are also meant towards increasing drug awareness as well as providing information regarding rehabilitation and treatment. Hence, we need to filter out tweets that promote or report drug abuse. We next describe the classification approach adopted to do so. Taking a cue from the works related to the automatic identification of prescription drug abuse tweets [40, 39, 27, 33], we used a supervised classification based approach to filter out drug abuse tweets from the set of tweets collected using keyword search.

Steps of tweets classification: The following steps are implemented to classify the tweets:

1. We use a corpus of a manually annotated dataset of $6,656$ tweets [27] to train a binary classifier.

2. We use a semantic sentence embedding approach, named Sent2Vec [32], to generate feature vectors from the individual tweets. Compared to the n-gram based feature

generation approach proposed in [27] that generates a large set of features (around $11,000$), the feature set generated by Sent2Vec is much smaller (around 700).

3. We enrich the feature sets with certain other handcrafted features that have been used in literature to classify drug abuse tweets. These include the presence and count of (*a*) certain abuse-indicating keywords that may indicate frequent overdoses, co-ingestion, alternative motives and routes of drug admission [19, 20], and (*b*) keywords representing drug-related slangs and colloquial words[5].

Table 3.2 compares the 10-fold cross-validation accuracy (applied on the 6,656 annotated tweets) of 5 different classifiers in identifying the abuse and non-abuse tweets when either n-grams or sentence embeddings with hand-crafted features are used. We observe that SVM trained with a combination of sentence embeddings and hand-crafted features significantly outperforms the rest of the classifiers. Out of the 2.2 million tweets, the classifier labels around 0.5 million tweets (23% of total tweets), of $278,448$ unique users, as prescription drug abuse tweets.

The observations thus highlight the enormity of scale of the drug abusers active in the social network and the tremendous threat they can pose in spreading of the drug abuse menace. This also motivates the need to study the underlying network and user characteristics so as to have a deeper understanding of the needs to prevent the spread of such abuse tweets.

3.1.3 Limitations of the dataset

The dataset considered for this study contains information about the users, their followers and the users they follow, in addition to each user's prescription drug abuse tweets. The web-scraping API used to collect tweets only retrieves original tweets, i.e. it does not contain retweets. Hence the Twitter API[6] was used to collect retweets of drug abuse tweets.

[5]https://www.noslang.com/drugs/dictionary.php
[6]https://api.twitter.com/1.1/statuses/retweets/:id.json

A limitation of this API is that it only retrieves 100 most recent retweets of the Tweet. As a result, we couldn't retrieve complete re-tweet information of 66 tweets which had more than 100 retweets. Another limitation of this dataset is that it does not contain information about when a link in the PDAN was created, i.e. when a user followed someone else in the PDAN. As a result, the PDAN is considered as a static network and the dynamicity of the edges could not be considered.

3.1.4 Qualitative analysis of the tweets

We perform a qualitative analysis of the 2.2 million tweets that were collected in the dataset. Table 3.3 provides certain examples of both drug abuse and non-abuse tweets. Majority of the non-abuse tweets, (that contain drug abuse keywords, but do not promote or report drug abuse) can be related to spreading awareness and rehabilitation or related to reporting effectiveness or side-effects of drugs used during medical treatments. Finally, a small fraction of tweets contained lyrics of songs about drug addiction and recovery. Although the accuracy of our approach in identifying the drug abuse tweets is significantly high, however, a deeper analysis of the tweets that were misclassified (figure 3.3) reveals that they predominantly contain keywords indicating either a motivation or side effects of drug abuse, both of which are also important features for the detection of drug abuse tweets. For example *"Should I take Xanax for my anxiety? I need something for anxiety that won't make me high if that makes sense"* is a non-abuse tweet misclassified as abuse tweet, while tweets like *"The steroid shot has fully completely in. Dishwasher loaded, two heaps of laundry complete, garbage out, eating like a horse."* subtly hinting at drug abuse are misclassified as non-abuse tweets.

We next describe the process of creating the prescription drug abuse network (PDAN) and describe some of its characteristics that can be relevant to the spread of drug abuse.

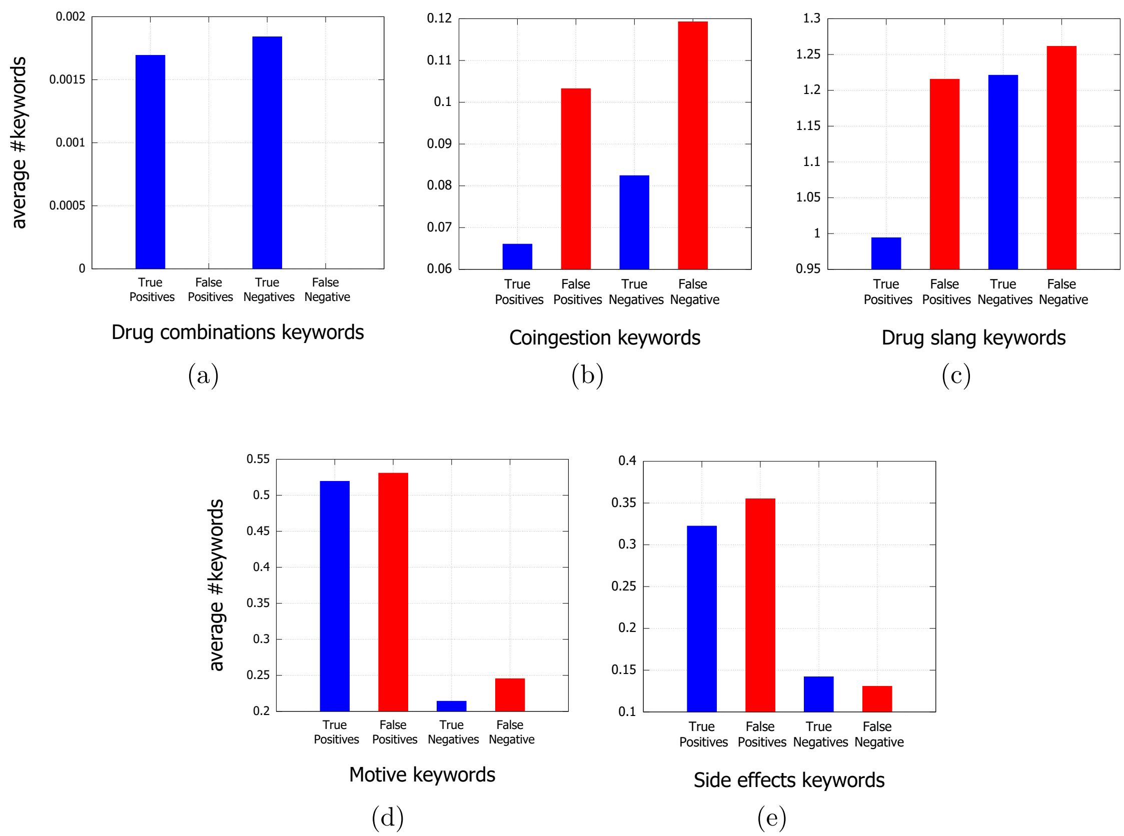

Fig. 3.3: Average values for each of the handcrafted features in the 4 categories of confusion matrix.

3.2 The Prescription Drug Abuse Network

In this section, we initially outline the steps of formation of the PDAN from the classified drug abuse tweets. Subsequently, to understand the role of the follower network in the spread of drug abuse, we observe the properties of PDAN and highlight how the network characteristics can be exploited in spreading drug abuse tweets to a significant number of users.

3.2.1 PDAN formation

We use the user information and their corresponding tweets to create the prescription drug abuse network (PDAN) of these users. To form the PDAN, we follow these subsequent

Non abuse tweets			Abuse tweets	
Awareness	**Rehabilitation**	**Treatment**	**Sale of drugs**	**Self reporting**
More than 17,000 opioid deaths in a year in US (2010). This includes prescription pain meds like **Percocet**, **Lortab**, **Vicodin**. It's not simply heroin any longer.	The Best Rehab Centers For **Vicodin** Addiction Treatment. Read Article: <*URL*>	Took a **Vicodin** & Zofran after a morning of treatment to control torment and queasiness and ran a 7 miler at 9:35 pace! #fcancer #livestrong	Buy **Percodan** Online now with no Prescription Or Members: buy **percodan** online now with no prescription or members... <*URL*>	I am absolutely abusing this **Vicodin** I got for my wisdom teeth.
Sarcasm	**Medical advice**	**Avoid drugs**	**Testimony**	**Asking for drugs**
90% of individuals in the medicinal field smoke weed, take Xanax, or adderall. Furthermore, the other 10% have lost their brains since they don't take it. #facts	@USER give him motrin. It works better than Tylenol.	This weird #anxietyattack can occur amid #pregnancy, and it's great to know so you don't take Advil to settle it <*URL*>	I was dependent on **Percocet** after my surgery. I had 6 refills in like, 2 & 1/2 weeks. my doctor took me off that so quickly. lmfao.	Does anyone have any **Vicodin** or anything they could sell me? Need it to keep me up over till tomorrow

Table 3.3: Example of the variety of tweets that match our keywords. The keywords (listed in table 3.1) that are used to search prescription drug abuse tweets are highlighted. All the tweets in this table are paraphrased to maintain the anonymity of the users.

steps.

1. We identify 278,448 unique users from the 0.5 million abuse tweets and retweets that we extracted and anonymize their identity for ethical concerns. The corresponding user mentions in the tweets were also suitably anonymized. All the tweets provided as examples in this thesis are paraphrased to maintain the anonymity of the users.

2. For each user, we use the Twitter API[7] to identify *follower-followee* relation between that user and the remaining unique users.

For precise understanding, a formal description of the prescription drug abuse network is as follows.

Let S denotes the set of drug abuse tweets as labeled by the classifier described in the previous section and $U = \{u_1, u_2, \ldots, u_n\}$ represents the set of users with at least one tweet or retweet in S. Thus if T_i denotes the set of drug abuse tweets and retweets made by u_i , then $T_i \cap S \neq \phi$. The PDAN is a directed graph represented as $G = \langle V, E \rangle$, where V is

[7] https://api.twitter.com/1.1/followers/list.json, https://api.twitter.com/1.1/friends/list.json

the set of nodes represented by the users in U and E is the set of directed edges between the node pairs. A directed edge is created from node j to i (denoted as e_{ij}) if user u_i is followed by u_j.

We next observe the properties of PDAN and investigate the possible support it can provide in spreading of drug abuse tweets.

Chapter 4

Characterizing the Prescription Drug Abuse Network

4.1 Characterizing PDAN

Network Property	**Value**
Number of Nodes	$278,448$
Number of links	$4,966,424$
Average (in/out) degree	19.05
In-degree slope	-1.51
Out-degree slope	-2.00
Number of Component	384
Giant Component Size	$259,623$
Clustering coefficient	0.1479
Reciprocity	0.6182

Table 4.1: Statistics of Prescription Drug Abuse Network.

Component	**Count**	**Percentage**
LSCC	226,254	86.77
IN	26,606	10.20
OUT	3,917	1.50
Tendrils	2,846	1.09
Disconnected	1,124	0.43

Table 4.2: Distribution of users of giant component in different Bowtie components.

We highlight some of the major network properties of PDAN in table 4.1 that would impact spreading in the network.

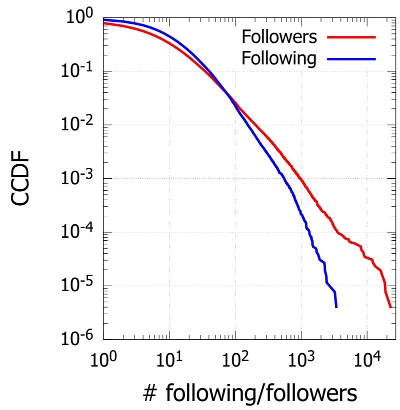

(a) CCDF of followings and followers.

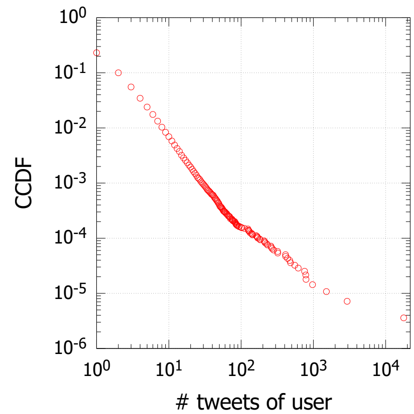

(b) CCDF of number of tweets of users in PDAN.

Fig. 4.1: (a) represents CCDF of the follower count and following count of the users and (b) represents CCDF of the number of tweets of users in the network.

4.1.1 Basic Statistics:

We observe that there exists a large prescription abuse network consisting of approximately 0.28 million unique users with 0.5 million links between them. The in and out degree of the nodes in PDAN follow power-law distributions with exponents 1.48 and 2.05, respectively. It may be noted that the exponent for the out-degree observed in the entire Twitter network was 2.276 [29]. On an average the users in PDAN have around 2,081 followers (including the followers who are not in PDAN) and follow around 680 user which indicates they are popular users. Figure 4.1(a), shows the complementary cumulative distribution function (CCDF) of the follower and following count of users in PDAN. As the graph indicates, a significant fraction of drug abusers have a very high number of follower and following count (42% users have more than 10 followers and 2%, i.e. 6,109 users, have more than 100 followers in PDAN), suggesting the possibility of certain users being able to directly reach a large fraction of the network and hence can play a dominant role in the spread of drug abuse.

4.1.2 Connectedness:

We observe that although there are 384 components in the network, the size of the giant component is 93.24% of the whole PDAN with $259,623$ nodes. But the second and the third largest connected components have 107 and 28 nodes respectively, and there exist several smaller components with an average size of around 3.

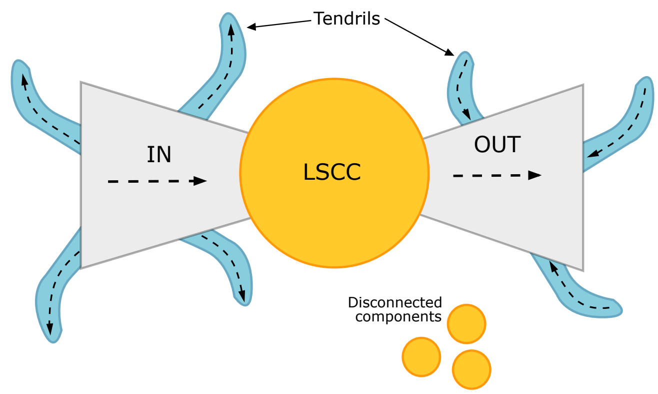

Fig. 4.2: Pictorial representation of the bowtie structure of PDAN.

Further, if we model the (giant component) network as a BowTie structure [6] (Figure 4.2) it is found that around 87% of the nodes (refer table 4.2) fall in the largest connected component (LSCC) or the core of the structure. The percentage of IN nodes, i.e. the ones that are following one or more nodes in the core is around 10% whereas the corresponding OUT percentage is around 1.5%. A minuscule of the nodes are either TENDRIL nodes (1%) (with no direct connectivity to the core nodes) or disconnected nodes (0.4%). The bow-tie connection often determines the resilience and robustness of the system. A large core implies that the network is resilient to targeted attacks, where bringing down a few nodes in the core will not prevent the information flow.

4.1.3 Clustering Coefficient:

Clustering coefficient is a measure of the degree to which nodes in a graph tend to cluster together. The undirected PDAN exhibits a high average clustering coefficient of around

0.15 that is significantly higher than observed in the actual Twitter network (0.096) [2]. This indicates the presence of small world property where friend of friend relations exist more commonly.

4.1.4 Reciprocity:

Reciprocity is the probability that if a user follows someone, she will be followed back by her. It is measured by the fraction of node pairs with reciprocating links. Our investigation further reveals the existence of a very high reciprocity (62% of total links) in the network. This value is significantly higher than the reciprocity value of around 22% observed in the twitter follower-followee network studied in [29].

All these statistics point towards a significant connectivity among the nodes along with high local clustering that are necessary for spreading and reaching out to a large fraction of nodes in the network. Further, the large and dense core of the system makes it robust to targeted removal of the nodes. Surprisingly, all these characteristics are consistent with several other networks, where behavioral spreading has been observed [8].

We next investigate the characteristics of the user's activities and their role in the PDAN.

4.2 Characterizing Users' activity

The activity of a user is determined by her activeness along with her role in the network. While activeness of a user depends on her engagement, the role a user can play depends on her positional importance. We initially provide measures to both activeness and positional importance and subsequently characterize the users based on both these parameters.

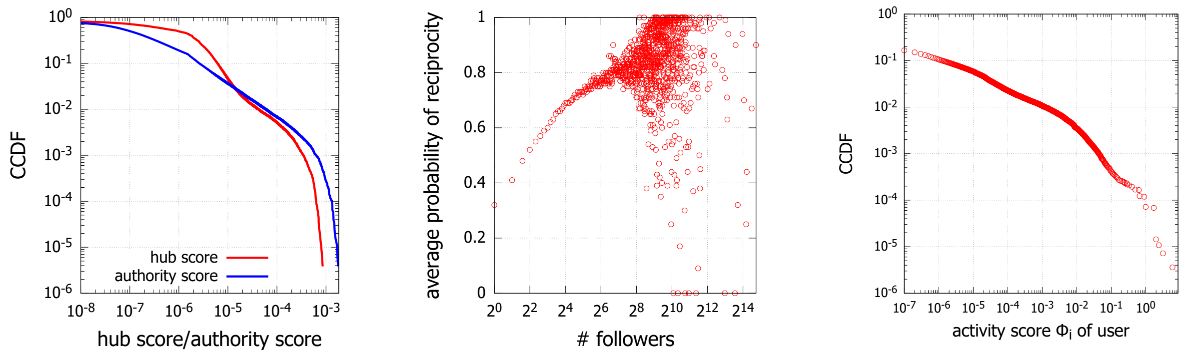

(a) Hub score and authority score for each user calculated using HITS algorithm.

(b) Average probability of reciprocity vs. number of followers.

(c) CCDF of user activity score in PDAN.

Fig. 4.3: (a) represents CCDF of the Authority and Hub score of the users and (b) represents the average probability of reciprocity as a function of the number of followers. (c) represents the CCDF of users activity score calculated as per equation 4.2.

4.2.1 Activeness and Positional Importance of Users

The activeness of a user provides a measure of her engagement in the PDAN whereas the positional importance reveals one's structural advantage in the spreading of drug abuse tweets. We next provide formal definitions of both these terms and the characteristics of the drug abusers with respect to these parameters.

Activeness of user

All users are not equally active on Twitter. Figure 4.1(b) shows the complementary cumulative distribution of the number of tweets of the users. The figure highlights that although most of the users are highly inactive (tweet count $\leq$ 2 about 90%), there exists a significant fraction of users with a very large number of tweets. Activeness can be measured based on user's tweeting behavior. We consider a user as active user if one *tweets more* with low *latency* (gap between two consecutive tweets). Thus the *activeness score* of user i, denoted as ϕ_i is defined as follows:

Users category	Number of users
Highly active users	$9,232$
Moderately active users	$55,226$
Inactive users	$213,990$

Table 4.3: Categorization user of PDAN based on user activity.

$$\phi_i = \begin{cases} \frac{|T_i|}{l_i}, & \text{if } |T_i| > 1 \\ 0, & \text{otherwise} \end{cases} \tag{4.1}$$

where, $T_i = \{t_1, t_2, \ldots, t_n\}$ is the set of sorted time stamp of the corresponding tweets of user i, $|T_i|$ is the total number of tweets by user i, and l_i is average latency between two consecutive tweets of user i that can be defined as follows:

$$l_i = \frac{1}{|T_i| - 1} \left(\sum_{k=1}^{|T_i|} (t_{k+1} - t_k) \right) = \frac{1}{|T_i| - 1} \left(t_{|T_i|} - t_1 \right) \tag{4.2}$$

The distribution of activity score as shown in figure 4.3(c) a heavy-tailed power law distribution. To categorize the users based on the activity score, we used Head/Tail breaks algorithm [23] to cluster the distribution into 2 parts. Users with $\phi_i > 4 \times 10^{-5}$ (i.e. the tail) were classified as *highly active* users. The remaining users were further classified into 2 categories, *moderately active* ($0 < \phi_i \leq 4 \times 10^{-5}$) and *inactive* ($\phi_i = 0$). Table 4.3 shows number of users belonging to each category according to activity score.

Positional Importance

The position of a user in the network can determine her reachability (ability to reach a large set of users through her tweets) as well as her accessibility (ability to receive tweets from a large number of users). To capture both these characteristics, we calculate the hub and authority score of each user in the network. Authority score provides a measure of the reach of a user while the hub score measures how accessible a node is from any other random node in the network. Thus, while authorities can act as good information spreaders, hubs can act as information collectors obtaining diverse information from different authorities. Hence hubs can play an important role in maintaining diversity in the discussion.

Hub \ Auth.	**Low**	**High**
Low	$248,863$	$8,209$
High	$10,550$	$10,826$

Table 4.4: Users' categorization based on hub and authority scores.

We used the HITS algorithm [28] to calculate authority and hub score of each user. Let $H(k)$ and $A(k)$ represent the hub and authority values of the user k, respectively. Figure 4.3(a) shows the CCDF of the hub and authority scores. We label the users as high authority user if its authority score is above 4×10^{-6} (obtained using Head/tail breaks algorithm [23]) and rest of the users are labeled as low authority users. Those users having hub score above 7×10^{-6} using the same algorithm are labeled as high hub user and rest of the users are labeled as low hub users.

We categorize the users into four role types based on authority and hub scores: ***a) information seeking*** – who have high hub and low authority scores, ***b) information sharing*** – who have high authority and low hub scores, ***c) leaders*** – who have high hub as well as high authority scores and ***d) fringe*** – who have low hub and low authority scores. In table 4.4, we observe that around 89% of the total users are fringe nodes who have few followers as well as very few followings. (as seen in table 4.6 and 4.7). On the other hand the total number of users in each of the remaining three categories i.e., information seeking, information sharing and leaders are only $3-4\%$. The users in these three remaining categories represent the influential section of users in the PDAN who have the capability to spread drug abuse tweets across a large section of the network. Thus it is necessary to investigate the contribution made by each of these user types in the spread of drug abuse tweets; hence we next correlate the activity score of the users with their hub and authority scores to identify the key players.

4.2.2 Characterizing highly active users

Initially, we take a closer look at the active users and observe their role types. The first row in table 4.5 shows the percentage of active users in each of the role categories. As can be observed, more than 90% of the highly active users are fringe nodes. This indicates that the fringe nodes are most actively involved in the spreading of drug abuse tweets.

Hops	**Info.Sharing**	**Leaders**	**Info. Seeking**	**Fringe**	**#users reached**	**% of network covered**
0	264	280	276	8,412	9,232	3.32%
1	4,980	8,064	4,404	59,560	77,008	27.66%
2	2,831	2,482	5,860	130,937	219,118	78.69%
3	58	0	10	24,722	243,908	87.60%
4	2	0	0	2,120	246,030	88.36%
5	0	0	0	178	246,208	88.43%

Table 4.5: Reach of highly active users at each hop. Hop 0 shows the distribution users highly active users by their role in PDAN followed by the distribution of users roles at each subsequent hop.

Hence, we further investigate the reachability of these active nodes at different hops. As shown in table 4.5, it is observed that in the first hop, the highly active nodes cover around 28% of the network. We find that 73% of these first hop users (3,090 info sharing, 3,152 leaders, 2,176 info. seeking and 47,823 fringe users in the first hop out of 77,008 users) are followers of active fringe. However, the active nodes cover a majority of the network (79%) only in their second hop indicating that with certain support from the first hop neighbors, the active users can effectively reach a large set of nodes. Thus this observation gears the need to focus on the followers of these highly active users and observe whether they can coordinate with the highly active users in the spreading process. We next highlight the observations with respect to these nodes.

User Role	Min.	$Q1$	Mean	$Q3$	Max.
Fringe	0	2	10.28	12	511
Info. seeking	0	2	10.17	13	214
Leaders	1	25	136.66	127	$27,243$
Info. sharing	1	23	122.30	112	$22,989$

Table 4.6: Distribution of in-degree for different user roles.

User Role	Min.	$Q1$	Mean	$Q3$	Max.
Fringe	0	4	12.76	16	358
Info. seeking	4	14	30.79	39	397
Leaders	3	42	118.09	131	$9,618$
Info. sharing	0	17	50.31	63	981

Table 4.7: Distribution of out-degree for different user roles.

Neighborhood of highly active users

We observe the neighborhood of the highly active users at each hop based on their hub/authority scores. As seen in table 4.5 while more than 90% highly active ($8,412$ users out of $9,232$) users are fringe nodes, however around 23% of the first hop neighbors of active nodes ($4,980$ info sharing, $8,064$ leaders and $4,404$ info. seeking users out of $77,008$ users) have either high hub or authority scores, i.e. non-fringe users, including more than 10% neighbors being leaders who have both high hub and authority scores. Because of the high hub score of their first hop neighbors, the active nodes reach around 79% of the total nodes by their second hop.

We next characterize the spreading process in the network and identify the role of the network structure and well as the users in the spread.

4.3 Characterizing Spread

We next focus our attention to the cascades of user engagement in Twitter. User engagement includes both generation of new tweets as well as re-tweets. We initially describe the experimental procedure to measure the cascade size and structural virality followed by the observations. Subsequently, the dynamics of spread are closely observed, keeping in mind the complex contagion phenomenon as observed in the spread of social behavior.

4.3.1 Measuring Cascades

We initially explain the method used to identify the cascades followed by the measures and the observations.

Dataset Preparation

The drug abuse tweets in the data set are initially sorted based on their time of generation. For each tweet, the corresponding creator is identified and is included as the initial node in the cascade graph $G = \langle V, E \rangle$. A user, v, following the initiator i is added as a node to the cascade graph if it has either re-tweeted or created a new drug abuse tweet within nine days of the appearance of the parent tweet in her timeline. The value of nine was chosen based on the study in [48], where it is shown that the mean of the attention decay time (time between peak attention and 75% of attention) of the tweets is around 217 hours. A directed link is created from node i to the follower v, indicating that engagement of node v has possibly been influenced by i. This process is further recursively repeated for the followers of the newly added users in G. Since prior to the first engagement of the user v, her last nine-day timeline can have drug abuse tweets from multiple users included in G, directed edges from all these users to v are also included, thus forming a graph structure. A user's particular tweet thus can be a member of several cascade graphs. We consider two graphs G_1 and G_2 as distinct if neither of the corresponding node sets V^{G_1} and V^{G_2} are subset of each other ($V^{G_1} \nsubseteq V^{G_2}$ and $V^{G_2} \nsubseteq V^{G_1}$). We include only the distinct graphs in our dataset.

Cascade Size and Structural Virality

One of the major characteristics of interest to the readers, about these cascade graphs, is the distribution of the graph sizes. The distribution of the cascade graph size provides an idea about how user engagement in drug abuse can be influenced by the drug abuse tweets. As shown in figure 4.4(a), the distribution of the cascade sizes follows a power law with an

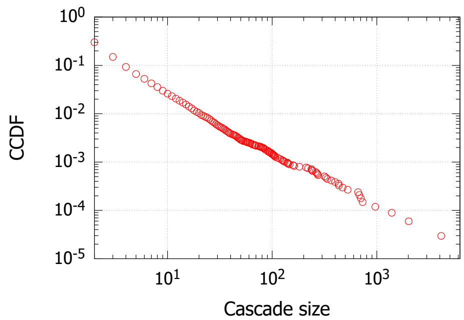

(a) CCDF of cascade size.

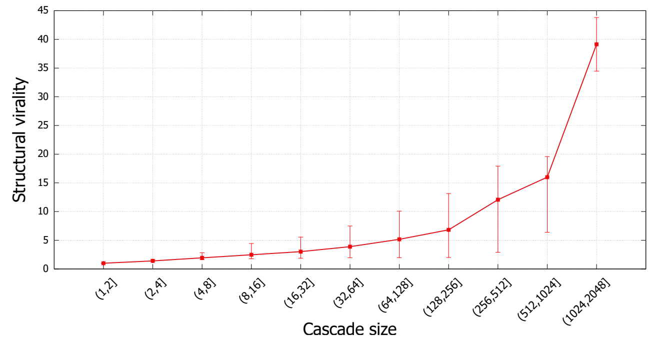

(b) Mean structural virality of cascades.

Fig. 4.4: (a) CCDF of cascade size and (b) shows the relation between cascade size and its mean structural virality (Wiener index).

exponent of 2.39. The maximum cascade size observed is 6, 230, which strongly indicates large cascades of user engagement may be formed due to the spread of drug abuse tweets.

We also observed the structural virality of these cascades, as defined in [17]. The structural virality measures cascade growth considering two possible extremes, one in which the virality is caused by a single node through a large broadcast and the other in which multiple users are engaged and a single node is responsible for only a part of the cascade. The structural virality has been measured using Wiener index that is given by the average shortest path length (d_{avg}) between any pair of nodes in the cascade graph. A lower value of d_{avg} (near to 1) indicates a hub-like structure where a single powerful node causes the entire cascade, whereas larger values indicate a chain like structure where multiple nodes are involved in the cascade. Figure 4.4(b) shows the distribution of the structural virality observed in PDAN. As can be observed, the structural virality increases rapidly with increasing cascade size. The structural virality is even significantly greater than 1 even for cascade size greater than 8. This indicates that multiple users are involved in the spreading of the drug abuse tweets and hence preventing spreading of drug abuse tweets cannot be accomplished by blocking only one or two users and thus can be quite challenging.

We next attempt to observe the dynamics of the spread of user engagement across the PDAN.

4.3.2 Influence and Dynamics of Spread

We next focus our attention to the role of content and influence of neighbors in the spreading of user engagement. We initially define the respective measures used followed by the observations made on our cascade data set.

Stickiness and Persistence

As rigorously observed in past studies on virality [41, 35], both the content and the influence of neighbors play dominant roles in virality. The probability of adoption of a content, spread through the neighbors of a user, is often measured by its *stickiness*. On the other hand, it has been observed that the probability of adoption of a content by a user is sometimes influenced by the number of exposures to that content through her neighbors (the complex contagion phenomenon) and has been measured using *persistence*. We avoid detailing the definitions of persistence in this paper, the details of which are present in [35]. However, we detail how our exposure curves are derived. In our case, we investigate how repeated exposure to drug abuse tweets influences the probability of adopting a similar engagement behavior. We subsequently investigate whether the probability varies for different drug names. A user is considered to be k-exposed if there are k users, whom the current user follows, have tweeted about drug abuse. We use an ordinal time estimate measure for deriving the exposure curve $p(k)$, whereby, the number of users ($I(k)$) who generates the first drug abuse tweet (indication of adoption) after being k-exposed but before being (k+1)-exposed is calculated and compared with the total number of k-exposed users ($E(k)$). The exposure curve is represented as $p(k) = \frac{I(k)}{E(k)}$. The stickiness is measured by the maximum value of $p(k)$ for all observed values of k and the persistence $F(p)$ is represented by the ratio of the area under the exposure curve and the minimum area of the rectangle covering the exposure curve entirely. $F(p)$ provides a measure of the rate of decay in the adoption probability with increasing number of exposures, after it has reached the peak. A value of $F(p)$ near to 1 indicates that repeated exposure to drug abuse tweets would be required before the user herself starts engaging, indicating the presence of a complex contagion phenomenon.

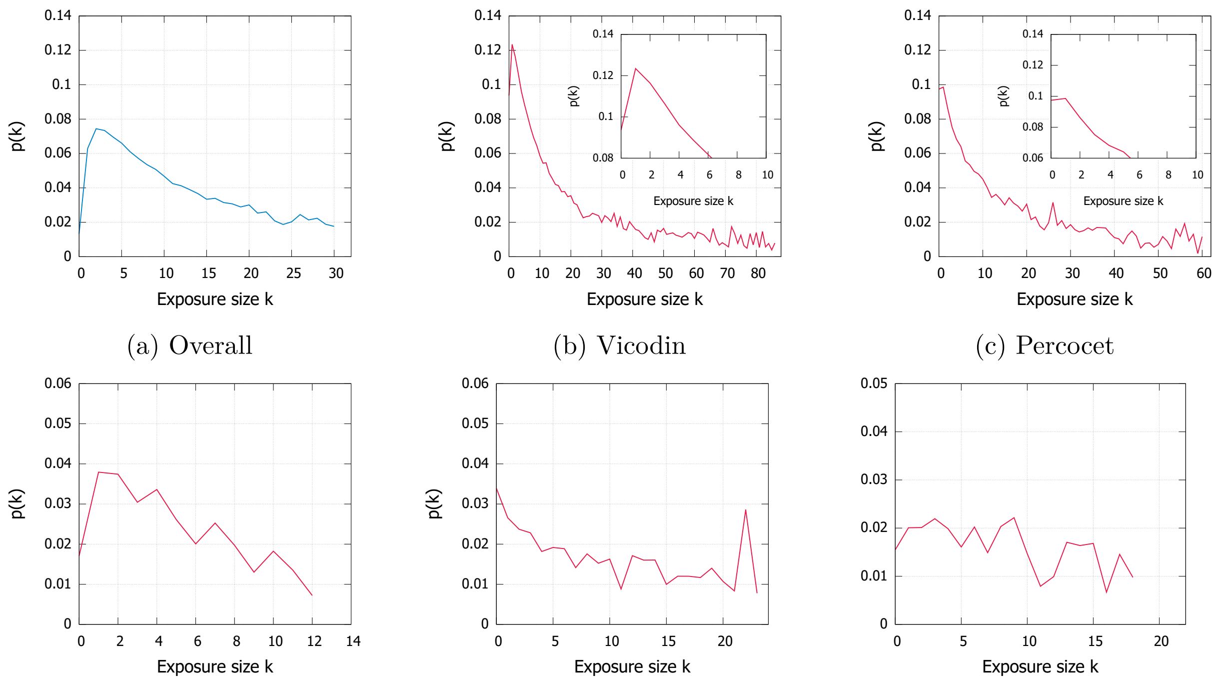

(a) Overall (b) Vicodin (c) Percocet

(d) Lortab (e) OxyContin (f) Dilaudid

Fig. 4.5: Average exposure curve for drug names. $p(k)$ is the fraction of PDAN users who tweet about a particular drug-name directly after their k^{th} exposure to it, given that they had not tweeted about it previously. The inset in figures (b) and (c) show the behavior of structural virality near the peaks at $k = 1$.

Observations

We obtained the value of $p(k)$ for each drug type present in our dataset. However, as the value of k increases, the amount of data required to calculate $p(k)$ decreases rapidly, making these observations error prone. To avoid this problem we consider values of k where the number of k-exposed users $E(k) > 500$. Figure 4.5 shows the average exposure curves for all the data and the five most popular drug-names (based on #exposures). We observe in figure 4.6(a) Vicodin and Percocet, that were found to be mentioned in significantly large number of tweets, have a relatively higher stickiness value of 0.12 and 0.10, respectively, compared to the other drugs. These values indicate that the maximum probability that a user engage herself with respect to these drugs are around 0.12 and 0.10, respectively. In both the cases peaks are found at $k = 1$, indicating that users mostly engage themselves about these tweets, after a single exposure only.

This high value of stickiness is observed for these abused drugs due to their high popularity on Twitter. In contrast Lortab, OxyContin and Dilaudid have a higher persistence of 0.61, 0.46 and 0.72 respectively as seen in figure 4.6(b), hinting that repeated exposures continue to have significant marginal effects.

To explain the exceptionally high stickiness values for Vicodin (figure 4.5(b)) and Percocet (figure 4.5(c)) at $k = 0$, we looked into the tweets containing these drug names. We observed that a sizable number of tweets containing Vicodin and Percocet are related to the sale of these drugs (we found $22,666$ and $18,094$ product sale tweets, respectively, for Vicodin and Percocet, compared to 46, 93 and 34 tweets for Lortab, OxyContin and Dilaudid, respectively). Since these tweets are generated independently, without being exposed, we see high values of $p(k)$ at $k = 0$ for Vicodin and Percocet. Further, since these drugs are popular among the drug abusers, repeated exposures to tweets related to these drugs do not lead to any significant effect on the engagement of the users. Users who are willing to discuss about these drugs, rapidly engage themselves after one or two exposures. Thus the persistence values of user engagement for these drugs are low. On the other hand, engagement for drugs, that are less popular over Twitter, shows high persistence. This could be possibly due to the fact that these drugs being less popular on Twitter, with increasing exposures the interest of the users about these drugs increases and hence the probability of engagement remains high with the number of exposures. Thus, there is a need to relook at the users who play a key role in the spread of these drugs. We investigate the same in next section.

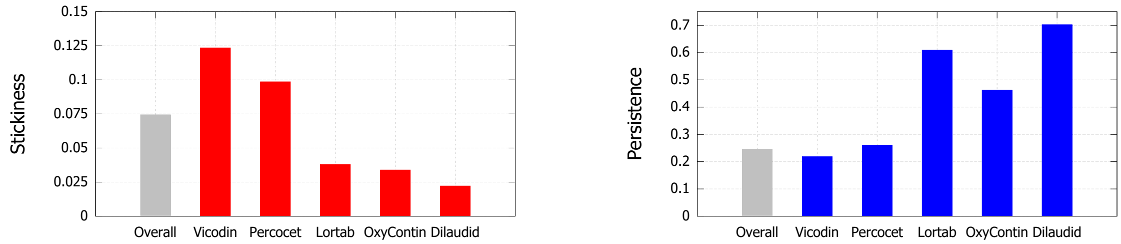

(a) Stickiness

(b) Persistence

Fig. 4.6: Stickiness and persistence score measures for drug names as defined in [35]

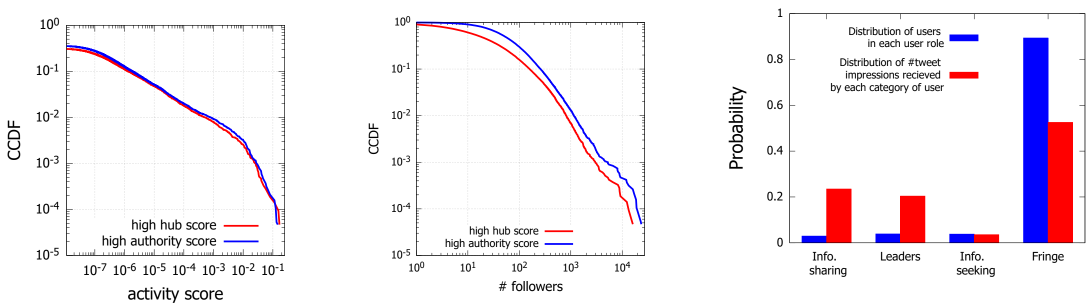

(a) CCDF of activity score of high hub and authority score users.

(b) CCDF of #followers of high hub and authority score users.

(c) Distribution of users in PDAN by their role.

Fig. 4.7: (a) represents CCDF of activity score of the high hub and authority score users and (b) represents CCDF of the #followers of the high hub and authority score users. (c) compares the distribution of number of users in each of the four categories defined above with the distribution of the number of impressions received by the users belonging to each of these category.

4.3.3 Key Players in Spreading

It has already been observed that the active nodes mostly comprise of fringe nodes that have low connectivity. On the other hand figures 4.7(a) and 4.7(b) show that most of highly connected users are not active. However, on investigating the probability that a tweet received by a random user is generated by a particular user role, the results reflect a strange paradox. It is observed that around 47% of the tweets (as seen in figure 4.7(c) info shares, leaders and info. seekers have 23.5%, 20.4% and 3.5% probability respectively) in the timeline of a random user is generated by a non-fringe node, even though they constitute only $10 - 12\%$ of the nodes in the network.

A closer look at table 4.6 reveals that in-degree of non-fringe nodes are 9 times higher (computed using weighted average) compared to in-degree of fringe nodes. So probability that a node follows a non-fringe node is comparable with the probability that the node follows a fringe node. Thus, since the fraction of fringe and non-fringe nodes followed by a random user is comparable, the proportion of active fringe and non-fringe nodes are also comparable, thereby generating contents in nearly equal proportions.

Cascade Property	**Cascades initiated by fringe nodes**			**Cascades initiated by non-fringe nodes**		
	Value	σ	Percent	Value	σ	Percent
#cascades initiated	148	-	42.41%	201	-	57.59%
Avg. cascade size	82.63	213.86	-	120.09	530.95	-
Avg. structural virality (d)	4.66	4.80	-	4.79	6.20	-
Max depth of cascade	128	-	-	156	-	-
Avg. depth of cascades	8.92	12.92	-	9.00	16.66	-
Avg. depth at which max. width was observed in cascade	5.53	10.45	-	4.53	11.48	-
Avg. #first hop nodes in cascade	8.57	22.72	-	10.23	17.10	-
Avg. #first hop non-fringe nodes	0.46	0.54	5.36%	3.11	4.52	30.38%
Avg. #first hop fringe nodes	8.11	22.87	94.64%	7.12	13.74	69.62%
Avg. #second hop nodes	4.74	8.77	-	5.63	15.66	-
Avg. #second hop nodes with non-fringe node as parent	3.20	8.44	67.48%	4.71	15.63	83.64%
Avg. #second hop nodes with fringe node as parent	1.54	3.94	32.52%	0.92	2.15	16.36%

Table 4.8: Comparison of the properties of cascades initiated by fringe and non-fringe nodes. The role of users play an important role in determining the properties of the cascade they belong to.

This observation gears the need to focus closely on the differences in cascade properties brought in by each type of user. For this experiment cascades of size ≥ 20 were considered. Table 4.8 compares the properties of these cascades initiated by fringe nodes and non-fringe nodes. It is evident that cascades initiated by non-fringe nodes have a greater size on average and more nodes (aggregates for all cascades) in its first and second hop. On investigating further, the presence of non-fringe nodes in the first-hop play an important role in inducting nodes in the second hop of the cascades. We can observe this phenomenon in both types of cascades, irrespective of the type of the initiator. This phenomenon is even more evident in cascades initiated by fringe nodes where only 5% of first-hop nodes belong to the non-fringe category but bring in 67% of the nodes in the second-hop. It is also observed that the maximum width of the cascades initiated by fringe nodes occurs at a much lower depth as compared to the ones generated by their counter parts. This further indicates that these cascades survive a few initial hops with the help of other fringe nodes only to peak later with the help of certain non-fringe ones.

We next attempt to uncover the prescription drug abusers. We use promoters identified by our approach to predict if a user is likely to tweet about prescription drug abuse. This could help in the early identification of potential drug abusers.

Chapter 5

Identifying the Promoters

5.1 Measures

In this section, we elaborate the procedure followed to identify the key masquerading promoters in the network. We use a modification of the Katz centrality based approach (for considering the edge weights) to rank the nodes based on their impact in spreading of drug abuse.

5.1.1 Measuring Node Impact

In this section, we provide a measure of the impact of a node in spreading drug abuse. We assume that the impact of a node is determined by two factors:

1. its potential to influence a large subset of its followers to drug abuse
2. its potential to influence a subset of its followers who themselves can have high impact in promoting drug abuse

We further assume that each node i will have a residual score(r_i) that measures the impact based on its individual activities (i.e. not dependent on its network properties) like the count and frequency of his drug abuse tweets. Based on these assumptions we use the PDAN to derive an expression for the impact score of the nodes. Since Katz centrality measure effectively captures our assumptions to derive the node centralities, we use a modification of the same for deriving the node impact scores.

If W represents the weight matrix with elements w_{ij} representing the weights of the ordered node pairs, then the impact, π_i of a node i would be given as

$$\pi_i = \alpha \sum_k w_{ik} \pi_k + r_i, \tag{5.1}$$

where $\alpha \in (0, 1)$ is a constant that controls the importance of the residual score r_i in the impact score; if $\alpha \to 0$, then $\pi_i \to r_i$ and hence the residual score gains more importance in determining the impact of the nodes. If Π and R represent vectors of the impact and the residual scores of the nodes respectively then using equation 5.1, Π can be expressed as

$$\begin{aligned} \Pi &= \alpha W \Pi + R \\ &= (I - \alpha W)^{-1} R \end{aligned} \tag{5.2}$$

where I is the identity matrix. Since α plays a critical role in the singularity of the matrix, there is need to determine a value of α for which the inverse of $(I - \alpha W)$ exists. Since $(I - \alpha W)$ is singular when $\det(I - \alpha W) = 0$, that implies $\det(W - \alpha^{-1} I) = 0$. Thus α^{-1} is an Eigen value of W. Hence for the inverse to exist, α^{-1} should be strictly greater than κ_1, the largest eigen value of W, i.e., $\alpha < \frac{1}{\kappa_1}$. Thus the node impact vector Π provides a ranking of the nodes based on their influence on other nodes and their role in propagating drug abuse. This method helps in identifying the top-k promoters in the network.

One of the major drawback of direct implementation of this approach is the large computation involved in calculating the inverse of the matrix. Hence we use a random walker based strategy to approximate the impact score of the nodes, which is described next.

5.1.2 Approximating Node Impact Score

We describe a lightweight random walker based strategy to derive a vector $\widehat{\Pi}$ that approximates the actual node impact vector Π derived in equation 5.2. Since our objective is to capture the node impact scores based on the influence of a node on its followers, we set the residual scores r_i of the nodes to 1. We ensure that strategy satisfies two major objectives:

1. the ranking of the nodes in Π must be preserved in $\widehat{\Pi}$
2. the strategy should be lightweight, i.e. it must gain significant accuracy with lower computational complexity.

We use a random walker technique to derive the node impact scores as mentioned in equation 5.1. The technique can be briefly described as follows: Each node in the PDAN initially sends D random walkers that traverse through an edge e_{ij} with a probability equal to the weight score (w_{ij}) of that link. It should be noted that w_{ij} is normalized between 0 and 1 and represents the personalized influence of a node i on its follower j. Thus, if node j sends D random walkers to node i with probability, w_{ij}, then node i receives an average of Dw_{ij} random walkers from node j. When a random walker from a source node j reaches node i after traversal of h^{th} hops, node i does any of the following two steps.

1. With a probability $\alpha^h w_{ki}$, it forwards the random walker to its neighbor node k, where α is same as defined above.
2. With a probability $1 - \sum_{l \in N(i)} \alpha^h w_{li}$, the random walker dies.

After all the random walker dies, the total number of random walker received by each node i is calculated, that provides an approximation of the node impact π_i. For a given network size and fixed α (based on the data set we consider the value of $\alpha = 0.125$), the complexity depends on the number of random walkers D sent by each node. However, with increasing values of D, $\widehat{\Pi}$ converges to Π. We avoid outlining the analytical proof of this convergence, however, we experimentally show the convergence of the Spearman rank

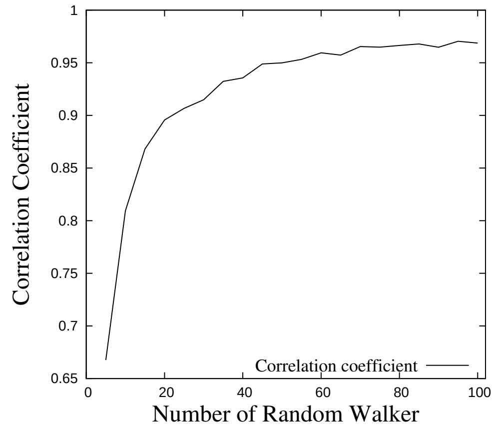

Fig. 5.1: Rank correlation between node impact measured by Random walker approach and value measured by equation 5.2.

correlation between $\widehat{\Pi}$ and Π with increasing values of D for a sample network of 1000 nodes (Figure 5.1). In the next section, we detail the experimental results highlighting the efficiency of the proposed approach.

5.2 Experiments and Results

In this section, we first validate the accuracy of our proposed approach in identifying the promoters and subsequently show that identification of these promoters helps in early detection of the potential drug abusers. To do so we introduce certain terminologies to categorize the users that are as follows,

1. ***Abusers***: Users present in the PDAN created with the prescription drug abuse tweet-feed up to time-stamp t. Hence the abusers have posted at least one drug abuse tweet.

2. ***Promoters***: Top 10% abusers, ranked by their node impact score π and with more than one tweet.

3. ***Potential abusers***: Followers of the current abusers who have not posted any drug abuse tweet (and hence not a part of PDAN) but posts a drug abuse tweet at some point in the future after time-stamp t.

4. ***Passive followers***: Followers of abusers (and not a part of PDAN) who will never post about prescription drug abuse and hence show no signs of being influenced by the promoters or abusers any time in future.

We next describe the validation dataset that we use for comparing the accuracy of the proposed approach with other approaches

5.2.1 Validation Dataset

To create the labeled dataset of the masqueraders, we consider only those users who are present in the PDAN. Subsequently, we use two important features for annotation. A drug abuser is labeled as a masquerading promoter if either:

1. The user has a very high follower count. A high follower count would imply that a user is influential and any drug abuse tweet made by the user would have a significant influence on his followers.

2. The drug abuse tweets of the user has received significant attention with high averages of favorite or retweet count per drug abuse tweet made by the user.

For each drug abuser, these values (follower count, average favorite count and average retweet count) were normalized by the corresponding maximum values. The highest among these three scores were considered for ranking the user. The top 10% of the abusers were labeled as masquerading promoters. It should be noted that none of these features have been used for calculating the influence score of the edges or in the impact score of the nodes. We next describe the baselines used for comparing our proposed approach and the measures used to evaluate the results.

5.2.2 Baselines

We have used the random walker technique described in section 5.1 to measure the node impact scores and subsequently identify the masquerading drug promoters. To compare the effectiveness of this technique we use other commonly used centrality measures like the degree centrality, weighted degree centrality, eigen vector centrality and PageRank to derive the node impact scores from the personalized influence scores. We use an additional node centrality measure, that we name as *TweetRank*, based on the count of the abusive tweets sent by the node. This baseline is important as it would justify the use of the proposed node impact score in identifying the key promoters.

5.2.3 Evaluation Measures

For all the experiments we use the precision and recall as the primary measures. These values as described in equations 5.3 and 5.4 respectively, are used to calculate other measures like F_1 score and AUC. Thus if *relevant* denotes the set of all masquerading promoters in the validation set and *retrieved* represents the set of promoters that are identified using our approach, then

$$precision = \frac{|retrieved \cap relevant|}{|retrieved|} \tag{5.3}$$

$$recall = \frac{|retrieved \cap relevant|}{|relevant|} \tag{5.4}$$

We also use the AUC values (ranging between 0 and 1) to evaluate the performance of the proposed approach. To evaluate the accuracy in identifying future potential abusers and distinguish them from less vulnerable users, we use the F_1 score of the 2 classes (future abusers and non-abusers) along with the overall accuracy.

Next we discuss the accuracy of our proposed promoter identification approach.

5.2.4 Promoter Identification

We use the validation dataset described above to derive the precision and recall, as defined in equations 5.3 and 5.4 respectively, and the same are used to calculate area under the curve (PR-AUC) for our proposed approach as well as the baselines.

Centralities	Accuracy	PR-AUC
Random Walker	**50.20**	**0.39**
PageRank	30.46	0.17
Eigen Vector	30.37	0.11
Weighted Degree	35.56	0.18
Degree	45.69	0.35
Tweet Rank	15.03	0.03

Table 5.1: Accuracy and AUC calculate for the PR plots of several centrality methods used to identify the promoters in PDAN.

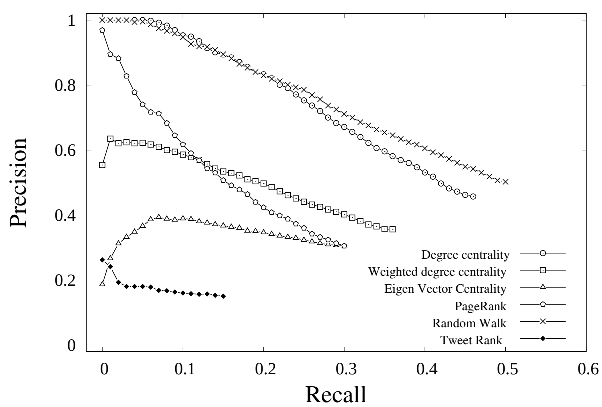

Fig. 5.2: Precision Recall curve for the five centrality measures.

The precision-recall (PR) curves and the corresponding AUC values are shown in figures 5.2 and table 5.1 respectively. To plot the precision-recall curve we initially consider the set of retrieved users as empty, i.e. $retrieved = \phi$ and the validation set, $relevent$, as the set of promoters. Iteratively we add an abuser with the highest centrality to the retrieved set and measure the precision and recall values. The values of precision and recall measured are then used to plot the PR-curve and calculate the AUC value.

As seen in table 5.1, our proposed strategy is most effective when compared to other baselines. Results indicate that the accuracy of the proposed approach is around 50% and the precision-recall area under the curve is around 0.4. Although these values are not very high, but they are comparable with the state-of-the-art top-k influencer identification approaches applied in other domains [42]. As PageRank technique lowers the effect of a follower with high follower count on the centrality value of a node, the node impact score of several influential promoters (who themselves are connected to other promoters) are getting reduced thereby lowering the accuracy values. Further the TweetRank measure that only considers the count of the tweets sent by an abuser as its centrality score yields the worst accuracy value. This shows that the key promoters are not the ones who advertise more, rather they use their social influence to increase the spread of drug abuse. This strongly motivates our approach of deriving the node impact based on the social influence exerted on the followers to identify the key promoters.

Next, we show that the impact score of the nodes can be used to predict future drug-abusers. We next outline a method of predicting potential abusers and evaluate its accuracy.

5.2.5 Identifying Potential Abusers

We term certain tweet users as potential abusers who are currently not abusing drugs but are vulnerable and may become drug-abusers in near future. To identify the potential abusers, we exploit the fact that a user following either certain influential drug promoters or considerable number of other drug abusers are more vulnerable to drug abuse.

We use a supervised learning approach that uses the impact score π of a drug abuser in the PDAN along with his follower information as features to identify potential abusers. The impact score of the nodes are calculated based only on the drug abuse tweets made over a period of 6 months. We train a binary classifier which in turn will distinguish potential abusers from the passive followers. We use the following followee information of a user as the feature set: the total number of followees, their total node impact score and their maximum and minimum influence scores. To train the classifier we use 2 sets of nodes in equal proportions. A node is labeled as a potential abuser if the user makes a drug abuse tweet for the first time in the next 6 months. We also include certain nodes who follow one or more drug abusers in the PDAN, included in the first 6 months of the training set but remain as passive followers for the entire duration of 5 years. The classifier is iteratively

Classifier	F_1 **Potential Abuser**					F_1 **Passive Follower**					**Accuracy**				
	6 m	12 m	18 m	24 m	30 m	6 m	12 m	18 m	24 m	30 m	6 m	12 m	18 m	24 m	30 m
Considering node impact															
NB	0.661	0.704	0.632	0.687	0.692	0.601	0.581	0.393	0.240	0.276	72.31	67.30	57.51	61.14	59.85
SVM	**0.729**	**0.751**	**0.707**	**0.715**	**0.720**	**0.696**	**0.689**	**0.570**	**0.644**	**0.685**	**71.30**	**72.33**	**65.14**	**68.35**	**70.39**
RF	0.596	0.627	0.611	0.677	0.659	0.688	0.409	0.523	0.657	0.611	64.80	54.23	57.14	66.76	63.64
LR	0.744	0.721	0.666	0.523	0.523	0.699	0.605	0.415	0.672	0.653	72.31	67.30	57.51	61.14	59.85
EN	0.734	0.738	0.670	0.720	0.741	0.685	0.680	0.574	0.524	0.697	71.13	71.19	62.83	64.70	72.09
Without considering node impact															
NB	0.649	0.660	0.665	0.664	0.662	0.096	0.052	0.066	0.069	0.062	49.45	49.98	50.64	50.63	50.27
SVM	0.667	0.667	0.667	0.667	0.667	0.000	0.000	0.000	0.000	0.003	50.02	50.02	50.01	50.01	50.08
RF	0.631	0.622	0.652	0.643	0.606	0.625	0.465	0.574	0.605	0.467	62.81	55.69	61.68	62.47	54.71
LR	0.514	0.446	0.445	0.512	0.508	0.295	0.386	0.444	0.326	0.315	42.46	41.75	44.46	43.41	42.73
EN	0.694	0.668	0.660	0.676	0.678	0.568	0.372	0.561	0.576	0.550	64.17	56.58	61.69	63.30	62.45

Table 5.2: Results of future predictions by training the model iteratively by adding 6 months data. (NB = Naive Bayes, RF = Random Forest, LR = Logistic Regression, EN = Ensemble)

trained with an additional 6 months of temporal data (i.e. we increment the training set with $6, 12, 18 \ldots, 30$ months of temporal data) and using the next 6 months dataset for validation.

Table 5.2 shows the accuracy of this prediction process. We used a variety of models for comparison. AdaBoost with decision tree as the base classifier was used as the ensemble model. We observe that using SVM as well as the ensemble model we can achieve an average accuracy of more than 70%.

To observe the importance of using the edge influence and node impact scores as features in this identification process, we also trained the classifiers without using any of these features. The classifier is trained based on the number of followees of a node along with the count of drug abuse tweets received by it. Observations shown in table 5.2 indicate a significant drop in the F_1 scores and accuracy values for all the classification models, with around 20% drop in several cases including SVM. This highlights the importance of the impact score of the abusers used in our approach in influencing the potential abusers towards drug abuse.

Chapter 6

Conclusion

This thesis provides a detailed analysis of the Twitter follower network involving around 0.28 million users along with their characteristics that are involved in the promotion of prescription drug abuse using the Twitter platform. The follower network of drug abusers that we term as Prescription Drug Abuse Network (PDAN) was rigorously analyzed to understand whether it provides an inherent support in the spreading of the tweets. It was observed that the users in the PDAN organize themselves into a heavy core structure with high local connectivity among themselves. Such a large connected core provides various alternate channel of communication among the users, thereby providing an inherent support for spreading of the tweets. We studied the spread of the drug abuse tweets, related to various drugs, that encourages other abusers towards discussion and also acts as social advertisements for promoting such abuses. It was observed that large cascades of user engagement were generated with a significant percentage of them being initiated and driven by users with low positional importance (with low count of followers as well as its followings), that we term as fringe nodes in the PDAN. Also the structural analysis (measurement of structural virality) of the cascades formed shows that they are mostly not generated by a few important nodes, but is a collective phenomenon involving both the important as well as the fringe nodes. The findings lead to a proposition that the spreading nature of cascades is inherently resilient to targeted elimination of a few key nodes - this is a matter of concern and necessitates deeper and more detailed study of the network by the research community in the immediate future.

We also proposed an approach for identifying drug promoters. We analyzed the influence of these abusers on their respective followers in PDAN. We used a network centrality based approach to derive an impact score of each of these abusers in the network that provided a measure of the significance of that abuser in influencing other followers towards drug abuse. Experimental results on a validation dataset indicated that the proposed approach outperformed all other baseline approaches and identified the promoters with more than 50% accuracy. We further showed that the influence and node impact score can also be used for early prediction of potential future drug abusers, i.e. ones who are not currently abusing drugs but are at severe risk. Experimental observations on a temporal data set indicated that by using the influence and impact scores of the followees of a node as features, suitable classification models can identify the potential drug abusers with more than 70% accuracy. However, these predictions can be possibly improved if the tweets of the users can be analyzed to observe their psychological behavior. Further, observing the dynamics of tweet spread in the social network can also provide interesting insights about the drug promoters as well as the vulnerable users.

References

[1] N. Alvaro, Y. Miyao, and N. Collier, "Twimed: Twitter and pubmed comparable corpus of drugs, diseases, symptoms, and their relations," *JMIR Public Health and Surveillance*, vol. 3, no. 2, 2017.

[2] S. Aparicio, J. Villazón-Terrazas, and G. Álvarez, "A model for scale-free networks: Application to twitter," *Entropy*, vol. 17, no. 8, pp. 5848–5867, 2015. [Online]. Available: http://www.mdpi.com/1099-4300/17/8/5848

[3] S. Aral and D. Walker, "Identifying Influential and Susceptible Members of Social Networks," *Science*, vol. 337, no. 6092, pp. 337–341, jul 2012. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/22722253 http://www.sciencemag.org/cgi/doi/10.1126/science.1215842

[4] E. Bakshy, J. M. Hofman, W. A. Mason, and D. J. Watts, "Everyone's an influencer," in *Proceedings of the fourth ACM international conference on Web search and data mining - WSDM '11*. New York, New York, USA: ACM Press, 2011, p. 65. [Online]. Available: http://portal.acm.org/citation.cfm?doid=1935826.1935845

[5] E. Bakshy, I. Rosenn, C. Marlow, and L. Adamic, "The role of social networks in information diffusion," in *Proceedings of the 21st international conference on World Wide Web - WWW '12*. New York, New York, USA: ACM Press, 2012, p. 519. [Online]. Available: http://dl.acm.org/citation.cfm?doid=2187836.2187907

[6] A. Broder, R. Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tomkins, and J. Wiener, "Graph structure in the web," *Computer networks*, vol. 33, no. 1-6, pp. 309–320, 2000.

[7] C. Buntain and J. Golbeck, "This is your twitter on drugs: Any questions?" in *Proceedings of the 24th International Conference on World Wide Web*. ACM, 2015, pp. 777–782.

[8] D. Centola, "The Spread of Behavior in an Online Social Network Experiment," *Science*, vol. 329, no. 5996, pp. 1194–1197, sep 2010. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/20813952 http://www.sciencemag.org/cgi/doi/10.1126/science.1185231

[9] J. Cheng, L. Adamic, P. A. Dow, J. M. Kleinberg, and J. Leskovec, "Can cascades be predicted?" in *Proceedings of the 23rd international conference on World wide web - WWW '14*. New York, New York, USA: ACM Press, 2014, pp. 925–936. [Online]. Available: http://dl.acm.org/citation.cfm?doid=2566486.2567997

[10] M. D. Choudhury and S. De, "Mental health discourse on reddit: Self-disclosure, social support, and anonymity," in *ICWSM*, 2014.

[11] M. D. Choudhury, M. Gamon, S. Counts, and E. Horvitz, "Predicting depression via social media." in *ICWSM*, E. Kiciman, N. B. Ellison, B. Hogan, P. Resnick, and I. Soboroff, Eds. The AAAI Press, 2013.

[12] M. D. Choudhury and E. Kiciman, "The language of social support in social media and its effect on suicidal ideation risk," 2017. [Online]. Available: https://aaai.org/ocs/index.php/ICWSM/ICWSM17/paper/view/15662

[13] N. A. Christakis and J. H. Fowler, "The Collective Dynamics of Smoking in a Large Social Network," *N Engl J Med*, vol. 21358, no. 22, pp. 2249–58, 2008. [Online]. Available: http://www.nejm.org/doi/pdf/10.1056/NEJMsa0706154

[14] M. De Choudhury, "You're happy, i'm happy: Diffusion of mood expression on twitter," in *Proceedings of HCI Korea*, ser. HCIK '15. South Korea: Hanbit Media, Inc., 2014, pp. 169–179. [Online]. Available: http://dl.acm.org/citation.cfm?id=2729485.2729511

[15] T. Ding, A. Roy, Z. Chen, Q. Zhu, and S. Pan, "Analyzing and retrieving illicit drug-related posts from social media," in *Bioinformatics and Biomedicine (BIBM), 2016 IEEE International Conference on*. IEEE, 2016, pp. 1555–1560.

[16] M. Duffy and E. Thorson, "Emerging trends in the new media landscape," pp. 93–116, 01 2009.

[17] S. Goel, A. Anderson, J. Hofman, and D. J. Watts, "The Structural Virality of Online Diffusion," *Management Science*, p. 150722112809007, jul 2015. [Online]. Available: http://pubsonline.informs.org/doi/10.1287/mnsc.2015.2158

[18] M. Gomez Rodriguez, K. Gummadi, and B. Schölkopf, "Quantifying information overload in social media and its impact on social contagions," in *Proceedings of the Eighth International Conference on Weblogs and Social Media*. AAAI Press, 2014, pp. 170–179.

[19] C. L. Hanson, S. H. Burton, C. Giraud-Carrier, J. H. West, M. D. Barnes, and B. Hansen, "Tweaking and tweeting: exploring twitter for nonmedical use of a psychostimulant drug (adderall) among college students," *Journal of medical Internet research*, vol. 15, no. 4, 2013.

[20] C. L. Hanson, B. Cannon, S. Burton, and C. Giraud-Carrier, "An exploration of social circles and prescription drug abuse through twitter," *Journal of medical Internet research*, vol. 15, no. 9, 2013.

[21] J. Henrique, "Get old tweets-python computer api," https://github.com/Jefferson-Henrique/GetOldTweets-python, 2017.

[22] T.-A. Hoang and E.-P. Lim, "Tracking Virality and Susceptibility in Social Media," in *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management - CIKM '16*. New York, New York, USA: ACM Press, 2016, pp. 1059–1068. [Online]. Available: http://dl.acm.org/citation.cfm?doid=2983323.2983800

[23] B. Jiang, "Head/tail breaks: A new classification scheme for data with a heavy-tailed distribution," *The Professional Geographer*, vol. 65, no. 3, pp. 482–494, 2013.

[24] J. Kalyanam, T. Katsuki, G. R. Lanckriet, and T. K. Mackey, "Exploring trends of nonmedical use of prescription drugs and polydrug abuse in the twittersphere using unsupervised machine learning," *Addictive behaviors*, vol. 65, pp. 289–295, 2017.

[25] J. Kalyanam and T. K. Mackey, "A review of digital surveillance methods and approaches to combat prescription drug abuse," *Current Addiction Reports*, vol. 4, no. 4, pp. 397–409, Dec 2017. [Online]. Available: https://doi.org/10.1007/s40429-017-0169-4

[26] L. Katz, "A new status index derived from sociometric analysis," *Psychometrika*, vol. 18, no. 1, pp. 39–43, 1953.

[27] A. Klein, A. Sarker, M. Rouhizadeh, K. O'Connor, and G. Gonzalez, "Detecting personal medication intake in twitter: An annotated corpus and baseline classification system," *BioNLP 2017*, pp. 136–142, 2017.

[28] J. M. Kleinberg, "Authoritative sources in a hyperlinked environment," *Journal of the ACM (JACM)*, vol. 46, no. 5, pp. 604–632, 1999.

[29] H. Kwak, C. Lee, H. Park, and S. Moon, "What is twitter, a social network or a news media?" in *Proceedings of the 19th international conference on World wide web*. ACM, 2010, pp. 591–600.

[30] T. K. Mackey, J. Kalyanam, T. Katsuki, and G. Lanckriet, "Twitter-based detection of illegal online sale of prescription opioid," *American Journal of Public Health*, vol. 107, no. 12, pp. 1910–1915, 2017, pMID: 29048960. [Online]. Available: https://doi.org/10.2105/AJPH.2017.303994

[31] H.-W. Meng, S. Kath, D. Li, and Q. C. Nguyen, "National substance use patterns on twitter," *PLOS ONE*, vol. 12, no. 11, pp. 1–15, 11 2017. [Online]. Available: https://doi.org/10.1371/journal.pone.0187691

[32] M. Pagliardini, P. Gupta, and M. Jaggi, "Unsupervised Learning of Sentence Embeddings using Compositional n-Gram Features," *arXiv*, 2017.

[33] N. Phan, S. A. Chun, M. Bhole, and J. Geller, "Enabling real-time drug abuse detection in tweets," in *Data Engineering (ICDE), 2017 IEEE 33rd International Conference on*. IEEE, 2017, pp. 1510–1514.

[34] D. M. Romero, W. Galuba, S. Asur, and B. A. Huberman, "Influence and Passivity in Social Media." Springer, Berlin, Heidelberg, 2011, pp. 18–33. [Online]. Available: http://link.springer.com/10.1007/978-3-642-23808-6_2

[35] D. M. Romero, B. Meeder, and J. Kleinberg, "Differences in the mechanics of information diffusion across topics," in *Proceedings of the 20th international conference on World wide web - WWW '11*. New York, New York, USA: ACM Press, 2011, p. 695. [Online]. Available: http://portal.acm.org/citation.cfm?doid=1963405.1963503

[36] S. Rubya and S. Yarosh, "Video-Mediated Peer Support in an Online Community for Recovery from Substance Use Disorders," in *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing - CSCW '17*. New York, New York, USA: ACM Press, 2017, pp. 1454–1469. [Online]. Available: http://dl.acm.org/citation.cfm?doid=2998181.2998246

[37] K. Rudra, A. Chakraborty, M. Sethi, S. Das, N. Ganguly, and S. Ghosh, "#FewThingsAboutIdioms: Understanding Idioms and Its Users in the Twitter Online Social Network." Springer, Cham, 2015, pp. 108–121. [Online]. Available: http://link.springer.com/10.1007/978-3-319-18038-0_9

[38] D. Saez-Trumper, G. Comarela, V. Almeida, R. Baeza-Yates, and F. Benevenuto, "Finding trendsetters in information networks," in *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '12*. New York, New York, USA: ACM Press, 2012, p. 1014. [Online]. Available: http://dl.acm.org/citation.cfm?doid=2339530.2339691

[39] A. Sarker and G. Gonzalez, "Data, tools and resources for mining social media drug chatter," *BioTxtM 2016*, p. 99, 2016.

[40] A. Sarker, K. OâĂŹConnor, R. Ginn, M. Scotch, K. Smith, D. Malone, and G. Gonzalez, "Social media mining for toxicovigilance: automatic monitoring of prescription medication abuse from twitter," *Drug safety*, vol. 39, no. 3, pp. 231–240, 2016.

[41] D. A. Sprague and T. House, "Evidence for complex contagion models of social contagion from observational data," *PLOS ONE*, vol. 12, no. 7, p. e0180802, jul 2017. [Online]. Available: http://dx.plos.org/10.1371/journal.pone.0180802

[42] K. Subbian, C. C. Aggarwal, and J. Srivastava, "Querying and tracking influencers in social streams," in *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining*. ACM, 2016, pp. 493–502.

[43] G. Szabo and B. A. Huberman, "Predicting the popularity of online content," *Communications of the ACM*, vol. 53, no. 8, p. 80, aug 2010. [Online]. Available: http://portal.acm.org/citation.cfm?doid=1787234.1787254

[44] A. Tamersoy, D. H. Chau, and M. De Choudhury, "Analysis of smoking and drinking relapse in an online community," in *Proceedings of the 2017 International Conference*

on Digital Health, ser. DH '17. New York, NY, USA: ACM, 2017, pp. 33–42. [Online]. Available: http://doi.acm.org/10.1145/3079452.3079463

[45] C. Tan, L. Lee, and B. Pang, "The effect of wording on message propagation: Topic- and author-controlled natural experiments on twitter," in *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, 2014, pp. 175–185. [Online]. Available: http://www.aclweb.org/anthology/P14-1017

[46] L. Weng, F. Menczer, and Y.-Y. Ahn, "Virality Prediction and Community Structure in Social Networks," *Scientific Reports*, vol. 3, no. 1, p. 2522, dec 2013. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/23982106 http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC3755286 http://www.nature.com/articles/srep02522

[47] S. Wu, J. M. Hofman, W. A. Mason, and D. J. Watts, "Who says what to whom on twitter," in *Proceedings of the 20th international conference on World wide web - WWW '11*. New York, New York, USA: ACM Press, 2011, p. 705. [Online]. Available: http://portal.acm.org/citation.cfm?doid=1963405.1963504

[48] S. Wu, C. Tan, J. M. Kleinberg, and M. W. Macy, "Does bad news go away faster?" in *Proceedings of the Fifth International Conference on Weblogs and Social Media, Barcelona, Catalonia, Spain, July 17-21, 2011*, 2011. [Online]. Available: http://www.aaai.org/ocs/index.php/ICWSM/ICWSM11/paper/view/2877