

A Large Scale Study of the Twitter Follower Network to Characterize the Spread of Prescription Drug Abuse Tweets

Ryan Sequeira, Avijit Gayen, Niloy Ganguly, Sourav Kumar Dandapat, Joydeep Chandra

Abstract—In this paper, we perform a large-scale study of the Twitter follower network, involving around 0.42 million users who justify drug abuse, to characterize the spreading of drug abuse tweets across the network. Our observations reveal the existence of a very large giant component involving 99% of these users with dense local connectivity that facilitates the spreading of such messages. We further identify active cascades over the network and observe that cascades of drug abuse tweets get spread over a long distance through the engagement of several closely connected groups of users. 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, thus potentially complicating the process of arresting such cascades. Further, we discovered that the engagement of the users with respect to certain drugs like Vicodin, Percocet and OxyContin, that were observed to be most mentioned in Twitter, is instantaneous. On the other hand for drugs like Lortab, that found lesser mentions, the engagement probability becomes high with increasing exposure to such tweets, thereby indicating that drug abusers engaged on Twitter remain vulnerable to adopting newer drugs, aggravating the problem further.

Index Terms—Social computing, Twitter, Information retrieval, Biomedical informatics

I. INTRODUCTION

THE enormous popularity of social media like Twitter makes it suitable as an advertisement platform for promoting drug-abuse. Keeping in view the spread and impact of drug-abuse (there is an estimated count of 190,000 premature drug-related deaths¹), and the limitations of traditional “*prescription drug monitoring programs*” (PDMPs) in understating the severity of this issue [1] there is a need to have a deeper understanding of how social media is playing a role in promoting the 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 [2] and the promotion of drug-abuse [3]. Possible surveillance strategies for identifying such retailers and drug-abusers have also been well explored [4]. However, an important aspect that needs to be carefully investigated is the networked effect of Twitter that may amplify the spread of drug-abuse tweets among users. Preliminary studies of Twitter users in [5] reveal the existence

of drug-related social circles (densely connected neighbor set of a user) around certain active users who tweet frequently, mentioning the different effects that specific drugs produce. Although it is not completely clear whether such active circles can influence non drug-abusers towards abuse, however, from concepts like the “*uses and gratification theory*” and “*Health Communication Media Choice*” (HCMC) model [6], it can be argued that such discussions that glorify drug-abuse are likely to engage users who justify drug-abuse in further deliberations to satisfy their communication needs. Consequently, cascades of user engagement, if formed through such discussions, would volume up as social advertisements that would downplay the ill effects of drug-abuse and may influence vulnerable users towards such practice. In this paper, we discover and characterize such cascades over the Twitter follower network, where users participate in discussions that treat drug-abuse positively.

To identify these cascades, we propose a technique that uncovers around 0.42 million unique users who were engaged in either self-reporting or promoting prescription drug-abuse through tweets. The enormity of this number reflects the huge role being played by social media in promoting prescription drug-abuse on Twitter. However, the dynamics of these cascades are mainly driven by the users’ engagement behavior, as well as the underlying structure of the follower network. Hence in our study, we follow a principled approach by first investigating the underlying follower network of these unique users and subsequently the characteristics of these users in terms of their engagement behavior and positional importance in the network, before finally studying the key features of the cascades. The *follower* relation among these users is used to create a network with directed edges. Investigation reveals that this network is almost entirely connected with around 91.37% of the nodes being in the largest strongly connected component. We also observe very high reciprocity of the links in the network. 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 assessed the activeness of the users and found that around one-fourth of the users are active. Moreover, the reach of these active nodes is considerable, as they cover a substantial part of the network.

From the detailed study of the nature of the network and its participating users, it is clear that a structure susceptible to facilitate the formation of cascades exist. We discover

¹https://www.cdc.gov/nchs/data/health_policy/monthly-drug-overdose-death-estimates.pdf

around 50,409 cascades, some of them with sizes reaching to thousands and extending over several hops. The network study of such cascades reveals high structural virality, i.e. the cascades are not driven by a single *important* node, which in turn makes them difficult to control. Further, the network among the nodes participating in each cascade exhibit significantly high clustering coefficient and reciprocity in their follower relationships. All these observations provide a strong evidence that drug-abuse tweets traverse through several groups of closely connected users. Moreover, the study of the characteristics of these users reveals equally important contributions, irrespective of their position and activeness, in the spreading process. The phenomenon indicates that organic collaboration among a large set of nodes is largely responsible for the emergence of a cascade; thus it may be difficult to control the cascades by eliminating a few targeted nodes. Finally, guided by the metrics associated with social contagion processes [7] (discussed in detail later), we find that users engage (through tweets or retweets) differently when exposed to drug-abuse tweets with different drug names. On exposure to tweets with highly mentioned drug names like Vicodin, Percocet and OxyContin, the engagement is instantaneous, whereas for less popular drugs like Lortab the chances of engagement increase with increasing exposure to such mentions. This signifies that drug-abusers engaged on Twitter remain at risk of adopting newer drug types to which they are continually exposed through Twitter discussions.

The rest of the paper is organized as follows. In the next section, we highlight the related works. In section III, we describe the detailed dataset used in our work. In this section, we explain the data collection process as well as the classification method to identify the drug-abuse tweets. In section IV, we describe the follower network formation method and further detail its network characteristics. We subsequently discuss the user characteristics in section V. The details about the spreading pattern of the tweets across the follower network are outlined in section VI. Finally, we summarize our findings in section VII.

II. RELATED WORK

A plethora of recent work uses social media to gather information and in turn, provide solutions to various issues related to health. For example, social media has played a vital role in providing rich information for inferring mental health conditions (especially depression [8], mood instabilities [9], suicidal risks [10] and the effects of psychiatric medication [11]), as well as lifestyle-related conditions like overeating, alcoholism and smoking [12]. Technological approaches are being leveraged for addressing critical issues like the early prediction of such diseases [8], increased support and service engagement [13] and a decrease in the duration of untreated disorders [14]. These works provide a direction to the critical issues, concerning the psychological problems (including the drug-abuse problem), that require immediate attention.

Recently, prescription drug-abuse is receiving increasing attention due to its significant spread and the casualties involved [15], [16]. Majority of these work are directed towards

identifying content on social media that reveal drug-abuse behavior of the users. Such techniques include both supervised classification [17]–[20] as well as unsupervised methods [15], [21] for identifying drug-abuse tweets from the tweet stream. Complementary to the problem of drug-abuse detection, works based on identifying alternative opioid recovery treatments on social media also exist [22]. Such data can be leveraged to gain insights about the microscopic behavior of the corresponding users and their role in spreading drug-abuse tweets in the network, although only a few works have attempted to do so [2], [5]. One of the primary goals of this paper is to work towards these objectives.

In [5], 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 the abuse of specific drugs on Twitter) are likely to be surrounded by users who also participate in similar discussions, exhibiting a high content correlation among them. While this work provides preliminary insights about the tweeting behavior among these participating users, 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 content to understand human behavior has been a focus area across various topical domains. One of the earlier works on information flow on Twitter [23] showed the presence of a few elite users in Twitter who generate a majority of the content that is consumed by ordinary users. Several other works have also investigated the user characteristics and the role of influential users in information propagation across social networks [24], [25]. Subsequent empirical works have examined several other factors like the underlying network of the users [26]–[28], the user characteristics [29]–[31] the role of content [32], [33] and even the role of the underlying diffusion protocols [34] in such propagation. Selected works have investigated the effects of both users and content in information propagation [35], [36]. Since the dynamics of information propagation across networks vary with topics and content, this motivates the need to investigate the spreading behavior in the context of drug-abuse tweets by inspecting the cascades of user engagement, the role of the network and the user characteristics in the spreading process.

Spreading of various social behaviors has been investigated in several works, like the cessation of smoking [37], online sharing [38], and political controversies [7]. 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 through which such drug-abuse tweets spread and the role of the users that influence *spreading* of drug-abuse content, we believe this paper would contribute in filling this gap. In the next section, we describe

TABLE I: List of generic and brand names of prescription opioids medically used to treat pain².

Generic names	Brand names
oxycodone	OxyContin, Percodan, Percocet
hydrocodone	Vicodin, Lortab, Lorcet
diphenoxylate	Lomotil
morphine	Kadian, Avinza, MS Contin
codeine	-
fentanyl	Duragesic
propoxyphene	Darvon
hydromorphone	Dilaudid
meperidine	Demerol
methadone	-

the dataset used for this study.

III. DATASET

In this section, we provide a detailed description of the Twitter dataset along with the data collection methodology and the preprocessing techniques used. We subsequently describe the classification technique used to identify tweets that are promoting or reporting prescription drug-abuse. In line with the literature [5], the corresponding users engaged in such tweets are henceforth termed as *drug-abusers*. Based on the follower relation among these drug-abusers, a network is created. The steps followed to form the network is pictorially represented in figure 1.

A. Data Collection

The data collection steps can be briefly described as follows:

- 1) We prepared a set of drugs names (see table I) that have been marked and listed for abusive use in the past by the 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*”. This provided all the searchable tweets from January 2012 to July 2017 containing the drug names.
- 2) As only limited information about the tweets was being provided by the web-based crawler, we used the “*tweet id*” of the returned tweets to further query and extract the complete meta-data using the Twitter API³.
- 3) As retweets could not be retrieved using this web-based crawler, they were obtained using a different Twitter API⁴.
- 4) Finally, non-English tweets were identified using LangID⁵ and discarded.

Using this approach, we collected more than 2 million drug-related tweets. However, we observed that this collected tweet set included 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 applied several machine learning techniques to identify 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/show.json>

⁴<https://api.twitter.com/1.1/statuses/retweets/id.json>

⁵<https://github.com/saffsd/langid.py>

B. Classification of Drug-Abuse Tweets

TABLE II: Performance of binary classification of prescription drug-abuse tweets. 10-fold cross-validation is used to measure the F_1 score of the two classes (Drug-Abuse and No-Abuse) and the overall accuracy of the classifier.

Classifier	DA F_1	NA F_1	Accuracy
N-grams as features and handcrafted features			
Naive Bayes	0.752	0.701	72.87%
SVM	0.787	0.759	77.42%
Random Forest	0.842	0.794	82.12%
Logistic Regression	0.701	0.694	69.75%
AdaBoost	0.740	0.613	68.91%
XGBoost	0.776	0.694	74.13%
Bagging	0.770	0.752	76.14%
Voting	0.786	0.758	77.23%
Sentence embeddings and handcrafted features			
Naive Bayes	0.758	0.735	74.71%
SVM	0.857	0.846	85.16%
Random Forest	0.814	0.806	81.03%
Logistic Regression	0.811	0.803	80.70%
AdaBoost	0.782	0.770	77.61%
XGBoost	0.817	0.810	81.34%
Bagging	0.795	0.780	78.85%
Voting	0.843	0.832	83.83%
Deep learning with word embeddings			
LSTM	0.807	0.812	81.03%
RNN	0.819	0.820	81.99%
RCNN [39]	0.805	0.812	80.90%
TextCNN [40]	0.830	0.837	83.38%

One of the important requirements in identifying the drug-abusers is to classify the tweets based on whether they promote (or self-report) prescription drug-abuse or not. As a significant proportion of the tweets are meant towards increasing awareness against drug-abuse, we need to filter out tweets that promote or self-report drug-abuse from the rest of the tweets. Taking a cue from the works related to the automatic identification of prescription drug-abuse tweets [16], [17], [41], [42], we investigated several supervised classification based approaches to filter out drug-abuse tweets from the set of tweets collected using keyword search. We proceeded with text classification as follows:

Data Annotation: : We manually annotated 8,400 tweets, containing 14 different themes (see table III) that were identified through an exhaustive manual investigation made by 3 annotators on 20,000 tweets from the collected dataset. These 20,000 were selected randomly from the collected data. We subsequently, selected 600 unique tweets, randomly, from each theme for annotation to create a balanced annotated training set of 8,400 tweets. Each tweet was subsequently labeled as ‘Drug-Abuse (DA)’ or ‘Non-Abuse (NA)’ by 3 annotators and inter-annotator disagreements were solved by majority voting. The 14 themes from table III can be grouped into DA or NA categories with each group having 7 themes. Hence the resulting dataset used for measuring classification performance is balanced and has equal number of drug-abuse and non-abuse tweets (i.e. 4,200 tweets per category).

Feature Selection: : We use three different feature sets for binary classification, based on the classifier, as follows:

- 1) A combination of bi-gram vector and handcrafted features, proposed in [16].

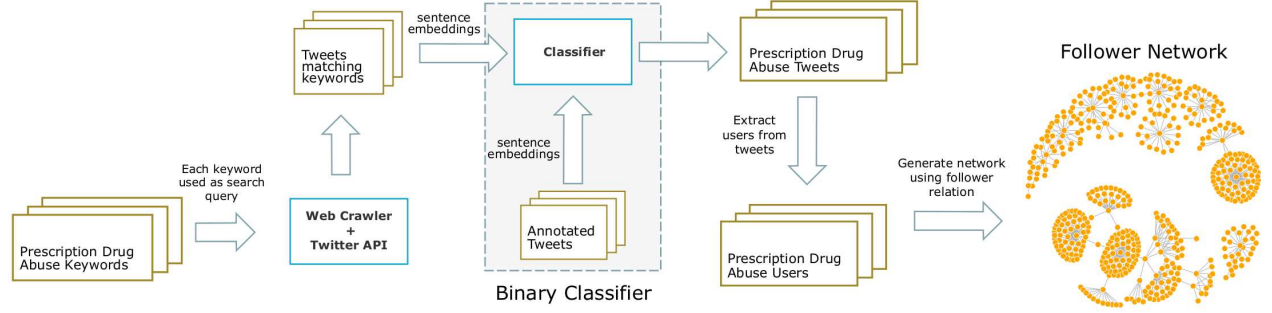


Fig. 1: Steps of follower network formation.

TABLE III: Example of the variety of tweets that match our keywords. The keywords (listed in table I) 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.

Drug-Abuse Examples		Non-Abuse Examples	
Category	Tweet	Category	Tweet
Addiction	I was an addict, a complete addict! I was consuming more than <NUMBER> mg of Morphine and OxyContin each day.	Metaphor / Sarcasm / Jokes	What you call Alabama Shakes we call that OxyContin withdrawal in the state of Ohio.
	I am officially addicted to Vicodin ... I need help.		My grandmother was telling about how much she money could earn by selling her Percocet on the streets. Imfao.
Co-ingestion	Washing Demerol down my throat with some vodka. I sense that I have officially failed at life.	Awareness	Left over Vicodin : Flush or Trash? Link to FDA for Safe Medication Disposal. <URL>. A useful homepage link for patients?
	Vicodin with Weed = a long sleep.		The CDC says <NUMBER> people died this year from prescription opioid overdose. # Percocet
Alternate modes of ingestion	I'm sitting on a large leather chair railing a ton of Dilaudid .	Rehabilitation	Vicodin Rehab in Florida, Florida Center for Recovery <URL>.
	I'm about to snort some of the OxyContin off my table #YOLO		Behind the Dependence on OxyContin and Transition to Heroin @ Trusted Heroin Rehab <URL>
Recreation	Have I at any point referenced how much I recreationally enjoy Vicodin ?	Pop-culture references	OxyContin , Xanax bars, Percocet and Lortab / Valiums, Morphine patches, ecstasy / It's all up for grabs
	Having a Vicodin party at my place, who wants to join? #insomnia		Don't shoot, you can't fight a viking on Vicodin . Can you?
Selling	I'm amassing all my remaining Vicodin and Percocet from my previous two medical procedures and selling them as soon as I get back to the burgh	News	Washington city devastated by OxyContin addiction sues Purdue Pharma — claiming drugmaker put profits over citizens <URL>
	Giving out Percocet , \$<NUMBER> a pill... Get your hands on them while they last!!		Heroin makes lethal comeback after OxyContin becomes more difficult to crush - Alaska Dispatch <URL>
Ingestion	One of the metrits is that I get to pop Vicodin like breath mints. So that's always a good thing.	Medical treatment	Consumed a Vicodin to get rid of the pain in my mouth only to throw out everything a few hours after my surgery.
	High as a kite after taking that Vicodin .		Surgery went well, and i'm happy that it's over. Its time to take some Vicodin !!
Illegal online pharmacies	Order High quality OxyContin Online, <NUMBER>% discount. No Prescription Required <URL>	Pain	Sadly, the Percocet isn't helping me with my knee pain. I may need to explore something different.
	Did someone say Percocet ? Buy Percocet Online <URL>		My knee is in great pain. I need to take a Percocet for the pain.

- 2) A combination of semantic sentence embedding, Sent2Vec [43], and handcrafted features.
- 3) Word embeddings using GloVe.

Compared to the n -gram based feature generation approach proposed in [16] that generates a large set of features (around 11,000), the feature set generated by Sent2Vec is much smaller (around 700). The handcrafted features used in literature to classify drug-abuse tweets include the presence and count of (a) specific abuse-indicating keywords that may indicate frequent overdoses, co-ingestion, alternative motives and routes of drug admission [5], [44], and (b) keywords rep-

resenting drug-related slang and colloquial words⁶. Each word in GloVe embedding was represented using a 100 dimension vector.

Classification: : We use 10-fold cross-validation to report classification results. In each iteration of the 10-fold, tweets in the training set and test set were sampled randomly, ensuring an equal proportion of drug-abuse and non-abuse tweets in both the sets. This was done using the “Stratified K-Folds cross-validator” library of scikit learn, which ensured that the percentage of samples from each class were preserved in

⁶<https://www.noslang.com/drugs/dictionary.php>

training and testing set during each fold. In each fold, the sampling for the test set and training set was done with 1 : 9 ratio, ensuring that the tweets from the test sets of previous iterations are not repeated in the latest test set.

We used several machine learning as well as deep neural network based methods for classification. The machine learning models included Naive Bayes, SVM, Random Forest, Logistic Regression as well as Ensemble learning techniques like AdaBoost, XGBoost, Bagging and Voting. Decision Stump was used as the base classifier for AdaBoost and Bagging, while SVM, Logistic Regression and Naive Bayes classifiers were used as the base classifiers in Voting ensemble where the majority of the label predicted by the three classifiers was considered as the label. These models were trained on two different feature sets as, mentioned in (1) and (2). The deep learning models used include LSTM, RNN, RCNN [39] as well as TextCNN [40] and were trained on word embeddings mentioned in (3).

While LSTM and RNN are commonly used deep learning techniques for text classification, RCNN tries to improve upon them by incorporating contextual information in the recurrent structure. In contrast to the recurrent deep learning architectures, TextCNN adapts CNN for sentence classification, making it relatively faster to train and requires little hyperparameter tuning. All the deep learning models were trained with GloVe embeddings as inputs, where each word was represented by a 100 dimension vector. The LSTM model was parameterized with a dropout of 20% and recurrent-dropout of 20%. We used bidirectional GRU as the recurrent layer in RNN model and its output was max-pooled and average-pooled. The concatenation of the max-pooled and average-pooled vectors was given as inputs to the dense layer for classification. The RCNN model was parameterized to generate 100 dimensional (left and right) context vectors. In the TextCNN model, the kernel-size of the 3 (parallel) convolution layers were 2, 3 and 5 respectively. Each of its convolution layers had 128 filters. The output of the three convolution layers were max-pooled, concatenated and given as input to the dense layer. Table IV gives a summary of the important parameters for these DL models.

TABLE IV: Description of the model parameters.

Parameter	Value (Remarks)
LSTM	
Dropout	0.2
Recurrent Dropout	0.2
RNN	
Recurrent Layer	Type = Bidirectional GRU (Outputs of this layer were max-pooled and average-pooled.)
Dense Layer	Input = Concatenation of Average-pool and Max-Pool layers
RCNN	
Left Context	dim = 100
Right Context	dim = 100
TextCNN	
# Parallel CNN layers	3
CNN Layer 1	Kernel size = 2 , Filters = 128
CNN Layer 2	Kernel size = 3 , Filters = 128
CNN Layer 3	Kernel size = 5 , Filters = 128

Observations: : Table II compares the 10-fold cross-validation accuracy (applied on the 8,400 annotated tweets) of different Machine Learning (ML) and Deep Learning (DL) classifiers in identifying the abuse and non-abuse tweets. While DL classifiers generally perform better than ML classifiers, we observe that SVM trained with a combination of sentence embeddings and hand-crafted features outperformed all the classifiers. Based on the classification performance of SVM, it was used for further classification of 2.2 million tweets.

Although the accuracy of our approach in identifying the drug-abuse tweets is significantly high, however, we investigated some of the reasons for misclassification. From figure 2(a) it is evident that the correctly classified drug-abuse (DA) tweets prominently contain drug-abuse related slang terms (like oxycotton, hillbilly, percs, etc.) and keywords hinting at co-ingestion (like alcohol, beer, etc.), while correctly classified non-abuse (NA) tweets as seen in figure 2(c) have relatively low slang terms and co-ingestion keywords but higher mentions of motive keywords (like surgery or stress) and side-effects (like insomnia or migraine). While it is difficult to determine the exact reason for misclassification, but based on the evidence it is highly likely that NA tweets might be misclassified as DA (figure 2(b)) due to the presence of slang terms and co-ingestion keywords. On the other hand NA tweets might be misclassified as DA tweets due the relatively higher presence of motive terms and keywords hinting at side-effects.

Out of the 2.2 million tweets, the classifier identified around 0.77 million tweets (36% of total tweets), of 420,502 unique users, as prescription drug-abuse tweets. We subsequently used the Botometer [45] API ⁷ to identify and remove the bot accounts. This service assigns each user a bot score corresponding to the likelihood of that account being a bot. The score is calculated based on a set of 1,150 features using account metadata, content, network and temporal information. Using a threshold score of 0.5 [45], [46], 185 users were labeled as bots and their corresponding tweets were discarded.

For ethical concerns, we follow a rigid anonymization mechanism based on guidelines mentioned in [47]. The tweet and user identities were replaced by virtual identifiers. The corresponding user mentions in the tweets were also replaced by the corresponding virtual identifier. The timestamps of the tweets were also suitably replaced by relative values. All the tweets provided as examples in this paper were paraphrased and the URLs were also replaced by placeholders.

The observations thus highlight the enormity of the scale of the drug-abusers active in the social network and the tremendous threat they can pose in spreading of the drug-abuse menace. However, there are a few limitations of the dataset that we outline next.

C. Limitations of the Dataset

The dataset considered for this study contains information about the users, their followers and the users they follow on Twitter, in addition to each user's prescription drug-abuse

⁷<https://botometer.iuni.iu.edu/>

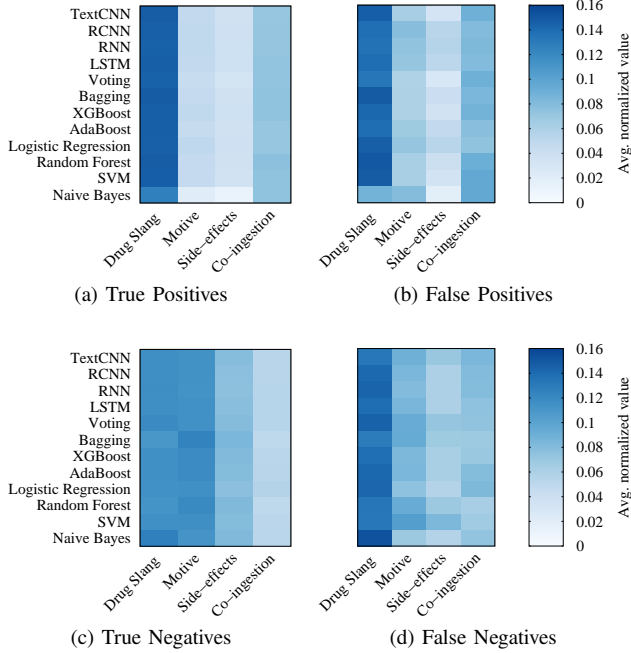


Fig. 2: Average normalized count of prominent keywords of (a) correctly classified DA tweets, (b) NA tweets classified as DA, (c) correctly classified NA tweets and (d) DA tweets classified as NA. ML classifiers trained on sentence embeddings and DL classifiers trained on word embeddings were considered.

tweets. The web-scraping API retrieves original tweets only, i.e., it does not contain retweets. Hence the Twitter API⁸ was used to collect retweets of drug-abuse tweets. 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 retweet information of 170 tweets which had more than 100 retweets. Another limitation of this dataset is that it does not contain information about the time when a user followed someone else. As a result, the network created from this dataset is considered as a static network and the dynamicity of the edges could not be considered.

D. Qualitative Analysis of the Tweets

We performed a qualitative analysis of the 2.2 million tweets that were collected in the dataset. Table III provides representative 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, news or be related to reporting effectiveness or side-effects of drugs used during medical treatments. Finally, a small fraction of tweets contained references to pop culture in the form of song lyrics about drug addiction and recovery or prescription drug-abuse references in movies or television shows.

As part of the study of the cascades of drug-abuse messages, we initially investigated the underlying follower network of the

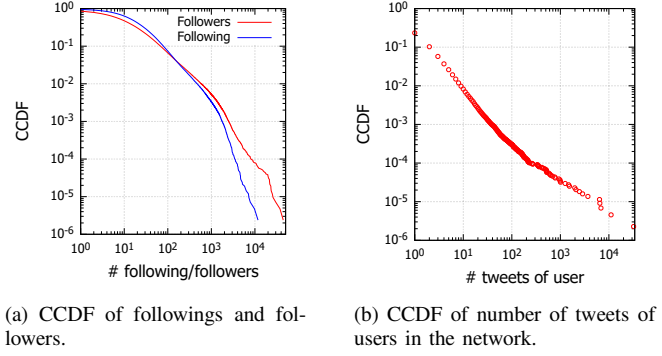


Fig. 3: (a) represents CCDF of the follower count and following count of the users and (b) represents CCDF of the number of tweets of users in the network.

drug-abusers along with their engagement characteristics, both of which are major drivers of these cascades. We next describe the process of creating the follower network and describe some of its characteristics that play major roles in the spread of the contents.

IV. THE FOLLOWER NETWORK

As the network provides the underlying framework for the spread of the drug-abuse tweets, we observe some of its essential properties and highlight their significance. However, we first briefly outline the steps of the formation of the network.

A. Network Formation

We used the user information and their corresponding tweets (or retweets) to create the follower network of these users. For each of the 420,317 unique users (denoted as U) identified from the 0.77 million abuse tweets and retweets, we used the Twitter API^{9,10} to identify the *follower* relation between a user and the remaining unique users. The resulting network is a directed graph represented as $G = \langle V, E \rangle$, where V is 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 the network and investigate the possible support it can provide in spreading of drug-abuse tweets.

B. Characterizing the network

We highlight some of the significant network properties (summarized in table V) that would impact the spreading of drug-abuse tweets.

⁸<https://api.twitter.com/1.1/statuses/retweets/:id.json>

⁹<https://api.twitter.com/1.1/followers/list.json>

¹⁰<https://api.twitter.com/1.1/friends/list.json>

TABLE V: Statistics of the follower network.

Network Property	Value
Number of nodes	420,317
Number of links	17,639,370
Average (in/out) degree	41.97
In-degree exponent	1.41
Out-degree exponent	1.38
Number of connected components	237
Giant component size	419,700
Clustering coefficient	0.1516
Reciprocity	0.6291

TABLE VI: Distribution of the users across different Bowtie components.

Component	Count	Percentage
LSCC	384,027	91.37
IN	29,578	7.04
OUT	4,467	1.06
Tendrils	1,224	0.29

Basic Statistics: We observed that there exists a large network consisting of approximately 0.42 million unique users, with 17 million links between them. The number of followers (in-degree) and followings (out-degree) of the nodes in this network follow power-law distributions with exponents 1.41 and 1.38, respectively. Figure 3(a), shows the complementary cumulative distribution (CCDF) of the number of followers and followings within the network. The in-degree exponent indicates that a significant fraction of drug-abusers have a very high number of followers (47% of users have more than 10 followers and 6.6%, i.e. 28,077 users, have more than 100 followers in the network), suggesting the possibility of a significant spread of the drug-abuse tweets if suitable connections exist among the nodes across the network. Consequently, we observed three major structural properties of this network — the connectedness among the nodes, the average clustering coefficient and the reciprocity of the links — all of which play an important role in the spread of tweets and can subsequently impact the user engagement.

Connectedness: To observe the connectedness of the network, we looked at the number of connected components, i.e., subgraphs where all the nodes within it are connected through a path and have no additional connections to any other nodes outside the supergraph. We observe that although there are 237 connected components in the network, however the largest (giant) component comprises of 99.85% of the total nodes (around 419,700 nodes), indicating that the network is almost entirely connected. The second and the third largest connected components have 62 and 20 nodes respectively, and there exist several smaller components with an average size of around 2. Further, if we model the (giant component) network as a BowTie structure [48] it is found that around 91% of the nodes (refer table VI) fall in the largest connected component (LSCC) or the core of the structure with much fewer nodes in the IN and OUT components. A large core, apart from being resilient to targeted attacks, also implies the possibility of faster diffusion of drug-abuse tweets over a significant fraction of the network [49].

Clustering Coefficient: This undirected network exhibits a high average clustering coefficient of around 0.15 that

is significantly higher than observed in the actual Twitter network (0.096) [50]. Although a high clustering coefficient has generally been considered as an impediment to large-scale diffusion across the networks [51], however, it needs to be investigated how the high clustering coefficient impacts the cascade properties in the drug-abuse networks.

Reciprocity: Our investigation further reveals the existence of very high reciprocity (63% of total links) in the network. This value is significantly higher than the reciprocity value of around 22% observed in the Twitter follower network studied in [52]. Previous studies have indicated that the existence of high reciprocal links not only affects the coverage of the spread of messages in social networks but also enhances the speed of the diffusion [53].

We later investigate the impact of these structural properties of the network on the cascades. However, these statistics point towards the existence of an underlying network platform that is amenable to the spread of the drug-abuse tweets through active engagement of a set of users. Hence we next attempt to characterize the users in the network by their engagement pattern.

V. CHARACTERIZING USERS IN THE NETWORK

In information diffusion, users playing dominant roles in the spreading processes have often been identified based on the importance of their contents, their positional significance in the network and their engagement time. In this section, we focus on the engagement time along with the positional significance to characterize the users and deal with the contents separately in the later section. While the engagement time or activeness of the users can contribute to the speed of diffusion, the positional importance can help in increasing the breadth or depth of the cascades [54]. We initially provide measures for both activeness and positional importance and subsequently characterize the users based on both these parameters.

A. Activeness of User

All users are not equally active in tweeting about prescription drug-abuse. Figure 3(b) shows the complementary cumulative distribution of the number of tweets of the users. The figure highlights that there exists a significant fraction of users with a vast number of tweets, with 4,286 users with more than 10 tweets and 134 users with more than 100 tweets. We consider a user as an active user if one *tweets more* with low *latency* (a gap between two consecutive tweets). The latency between two tweets is calculated based on the creation time of each tweet made available by the Twitter API. To avoid any confusion, we re-emphasize that Twitter does not provide the time when a user starts following someone. Thus the *activeness score* of user i , denoted as ϕ_i is defined as follows:

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


where, $T_i = \{t_1, t_2, \dots, t_n\}$ is the set of sorted timestamps of the corresponding tweets of user i , $|T_i|$ is the total number

TABLE VII: Categorizing user of the network based on their activity.

Users category	Number of users
Highly active users	4, 436
Moderately active users	95, 281
Inactive users	320, 600

of tweets by user i , and l_i is average latency between two consecutive tweets of user i that can be defined as follows:

$$l_i = \frac{1}{|T_i| - 1} \left(\sum_{k=1}^{|T_i|-1} (t_{k+1} - t_k) \right) = \frac{1}{|T_i| - 1} (t_{|T_i|} - t_1) \quad (2)$$



(a) CCDF of user activity score in the network.

(b) Hub score and authority score for users.

Fig. 4: (a) represents the CCDF of users' activity score as per equation 1, (b) represents CCDF of the Authority and Hub score of the users calculated using HITS algorithm.

The distribution of activity score as shown in figure 4(a) is a heavy-tailed power-law distribution. To categorize the users based on the activity score, we used the Head/Tail breaks algorithm [55] to cluster the distribution into 2 parts. Users with $\phi_i > 1.4 \times 10^{-3}$ (i.e. the tail) were classified as *highly active* users. The remaining users were further classified into 2 categories, *moderately active* ($0 < \phi_i \leq 1.4 \times 10^{-3}$) and *inactive* ($\phi_i = 0$). Table VII shows the number of users belonging to each category according to activity score.

B. Positional Importance

The position of a user in the network can determine her reachability (ability to reach a broad 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 simultaneously, we calculated the hub and authority score of each user in the network. Authority score of a node is high if it is followed by nodes with high hub score, whereas the hub score of a node would be high if it follows nodes with high authority scores. Thus, while authorities can act as good information spreaders because of their followers, hubs can act as information collectors obtaining diverse information from different authorities. We use the HITS algorithm [56] to calculate authority and hub score of each user.

Figure 4(b) shows the CCDF of the hub and authority scores. We labeled the users as high authority user if its authority score is above 6×10^{-6} (obtained using Head/tail

breaks algorithm [55]) and the rest of the users are labeled as low authority users. Those users having hub score above 4×10^{-6} using the same algorithm are labeled as high hub user, and the rest of the users are labeled as low hub users.

We categorized 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 IX, we observed that around 95% of the total users are fringe nodes who have few followers as well as very few followings. On the other hand, the total number of users in each of the remaining three categories is only 1 – 2%. The users in the three remaining categories represent the influential section of users in the network 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 explore their potential in spreading drug-abuse tweets.

C. Characterizing Highly Active Users

Initially, we took a closer look at the active users and observed their role types. The first row in table VIII shows the number of active users in each of the role categories. As can be observed, more than 96% of the highly active users are fringe nodes. This indicates that even though fringe nodes hold less positional importance but a significant fraction of the active engagements in generating drug-abuse contents is made by them. Thus to investigate the fraction of users getting exposed to the drug-abuse tweets generated by these active users, we further observed the reach of the active nodes at different hops.

As shown in table VIII, the active nodes by themselves can reach only 23% of the network in the first-hop but manage to reach 89% of the network in the second-hop. Thus, the reachability of the first-hop neighbors provides the active users the potential to reach the bulk of the nodes in the network. On taking a look at the properties of the first-hop neighbors, we find that around 13% of these nodes have either a high hub or authority score i.e. they are non-fringe users, including more than 9% leaders who have both high hub and authority scores. Roughly 0.21 million or 77% of the users in the second-hop follow at least one non-fringe user in the first-hop, highlighting the importance of the connectivity of non-fringe nodes. Interestingly, we also observed that 81% of the first-hop neighbors follow one or more active fringe nodes, indicating that even though individually the active fringe nodes are not structurally important but collectively they can reach a significant number of users in the first-hop.

Thus a major takeaway from these observations is that although most of the active users are positionally fringe, however, due to the positional importance of their first-hop neighbors, drug-abuse discussions initiated by these nodes can potentially reach a significant population of users in the network. We next observe the actual cascades and investigate

TABLE VIII: Reach of highly active users at each hop. Hop 0 shows the distribution of highly active users by their role in the network followed by the distribution of users roles at each subsequent hop.

Hops	Info.Sharing	Leaders	Info. Seeking	Fringe	#users reached	% of network covered
0	4	88	47	4,297	4,436	1.05%
1	374	8,828	2,741	79,985	96,364	22.93%
2	407	3,251	2,191	272,548	374,761	89.16%
3	12	0	0	37,155	411,928	98.00%
4	0	0	0	2,256	414,184	98.54%
5	0	0	0	130	414,314	98.57%

TABLE IX: Properties of different user roles.

User Role	#Users	Mean in degree	Mean out degree
Fringe	402,348	24.97	27.87
Info. seeking	4,979	72.48	133.06
Leaders	12,167	467.51	466.09
Info. sharing	823	1877.20	110.32

the role of the underlying network along with the key players and the tweet contents in the spreading process.

VI. CASCADES AND SPREAD

To measure the extent of the spread and influence of tweets, we investigate the cascades formed through drug-abuse discussions. We discovered the key players along with the pattern of user engagement responsible for the formation of the cascades. Here user engagement refers to both generation of new drug-abuse tweets as well as re-tweets by the users. Subsequently, the dynamics of spread with respect to the drug names are investigated, keeping in mind the complex contagion phenomenon that is typical to the spread of social behavior.

A. Measuring Cascades

We initially describe the experimental procedure to identify the cascades followed by the measures of various properties of the cascades.

Dataset Preparation: The cascades might be formed when users directly retweet a drug-abuse tweet or create a fresh content based on one that has appeared in their timeline. While retweets can directly be linked to a cascade, determining whether a new drug-abuse tweet has been made based on the previously received tweet (and hence should be a part of the cascade) is difficult. Possibilities exist that the user might have been influenced by certain external sources and not the drug-abuse tweet she has received before tweeting. However, ignoring such tweets entirely as not being part of the cascades can lead to a severe undermining of the veracity of the cascade problem. Hence to reduce the chances of such *coincidental errors*, following assertions were made: a) a tweet T was added to the cascade if at least one previous tweet from the cascade appeared on the user's timeline within a short period, Δ , prior to the sending of T and b) we consider only large cascades for analysis so as to ensure that a significant fraction of the spread does not suffer from such error. The value of Δ was chosen as nine days based on the study in [57], where it is shown that the mean of the attention decay time (time between peak attention and 75% of attention) of the tweets




Fig. 5: (a) represents CCDF of the cascade size (a) represents the relation between cascade size and mean structural virality (Wiener index) for each bin.

is around 217 hours. The process of identifying cascades is formally defined below.

The drug-abuse tweets in the dataset 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. 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 . By following this process, we extracted around 50,409 cascades of different lengths. We focus on the large cascades as they mainly reflect the threat posed by social media in the spread of drug-abuse tweets. To study the characteristics of the large cascades, we considered cascades of length ≥ 20 for our investigations. The choice of this value is not a principled one but is based on our observation that users with different characteristics play a consistent role across cascades beyond the length of 20. We next investigate the key structural properties of these large cascades.

Structural Properties of the Cascades: The distribution of the cascade sizes provides an idea about the scale of user engagement. As shown in figure 5(a), the distribution of the cascade sizes follows a power-law with an exponent of 2.06. The maximum cascade size observed is 23,164, which 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 [58]. The structural virality has been measured using the Wiener index that is given by the average shortest path length (d_{avg}) between any pair of nodes in the cascade

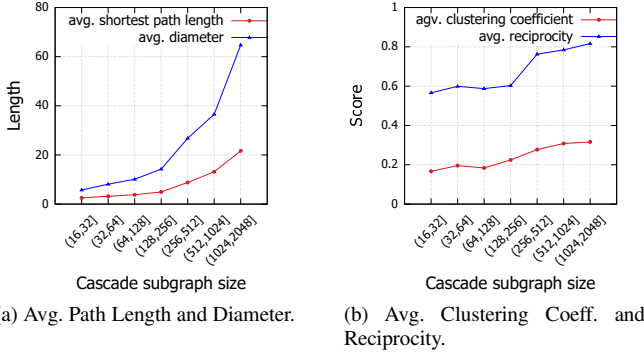


Fig. 6: Properties of cascade nodes in network subgraph. (a) shows the average shortest path and diameter of the cascade nodes in network subgraph and, (b) shows the average reciprocity and clustering coefficient.

graph. A lower value of d_{avg} (near to 2) indicates a hub-like structure where a single powerful node causes the entire cascade, whereas larger values indicate viral diffusion through branches involving multiple propagating nodes. Figure 5(b) shows the distribution of the structural virality observed in the follower network. As can be observed, the structural virality steadily increases with increasing cascade size. This indicates that multiple nodes play an important role in the cascades.

1) *Subgraph Properties of the Cascade Nodes:* To study the network structure of the nodes involved in the large cascades, for each cascade we constructed a network subgraph that comprises of the nodes of the corresponding cascade. As shown in figure 6(a), the average path length between the nodes in the subgraph vary from 2.51 to 21.66, whereas the corresponding average diameter ranges from 5.74 to 64.75 with respect to the subgraph sizes. These values indicate that these cascades reach far beyond the immediate neighbor nodes. It is also observed that the mean reciprocity and clustering coefficient of these subgraphs are significantly high (figure 6(b)). The mean reciprocity values with respect to the subgraph sizes vary between 0.57 and 0.81, whereas the mean clustering coefficients range between 0.17 and 0.31. A high average clustering coefficient and reciprocity of the cascade nodes provide strong evidence that the cascades spread along the follower chain through groups of closely connected nodes.

We next focus on the users to identify the key players involved in the spread of the drug-abuse tweets.

B. Key Players in Spreading

We investigated the key players in the cascades considering both the activeness as well as the positional importance of the users. Since a large fraction of the users in the network is largely inactive, we initially observed the activeness of the users in the cascade. We observed that for cascades involving more than 20 users, around 50 – 60% of the users are *one time engagers* who had contributed only one drug-abuse tweet but helped in keeping the cascade alive. This reveals the importance of these inactive users who contribute to increasing the length of the cascades.

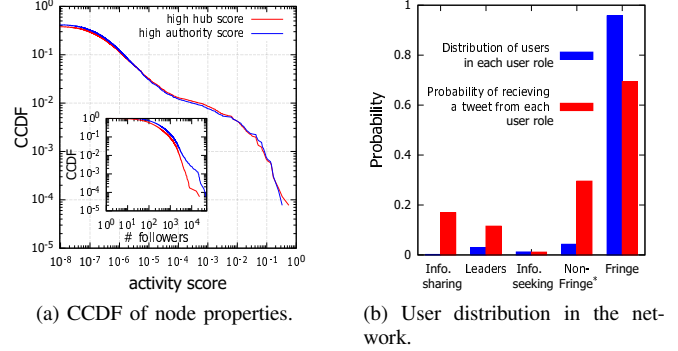


Fig. 7: (a) represents CCDF of activity score and the inset in (a) represents CCDF of the #followers of the high hub and authority score users. (b) gives the probability distribution of users in each category and the probability distribution of receiving a drug-abuse tweet from each user category. Non-Fringe category represents the combined probabilities of info. sharing, leaders and info. seeking users.

We next observe the role of the fringe and non-fringe nodes in the spreading process. Although it has been observed that the non-fringe active nodes are significantly less in number (refer table VIII) as compared to the fringe ones, however, it is observed that around 30% of the drug-abuse tweets (as seen in figure 7(b)) in the timeline of a random user are generated by a non-fringe node, even though they constitute only 5% of the nodes in the network. A closer look at table IX reveals that in-degree of non-fringe nodes are 17 times higher (computed using weighted average) compared to in-degree of fringe nodes. So the probability that a node follows a non-fringe node is comparable with the probability that the node follows a fringe node. Hence, as the fraction of fringe and non-fringe nodes followed by a random user is comparable, the contents generated by both these node types in the timeline of a random user is also comparable. This indicates that both the non-fringe as well as the fringe nodes are similarly responsible in the spread of these drug-abuse tweets.

C. Role of Neighbors in Spreading

We next focus our attention on the influence of neighbors in the spreading of user engagement. For this experiment, cascades of size > 20 were considered. Table X 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 with more nodes (aggregates for all cascades) in its first and second-hop. Irrespective of the type of the initiator, the presence of non-fringe nodes in the first-hop play an important role in inducing nodes in the second-hop of the cascades. This phenomenon is evident from cascades initiated by fringe nodes where only 5% of first-hop nodes belong to the non-fringe category but bring in 45% of the nodes in the second-hop. It is also observed that the maximum width of the cascades initiated by fringe nodes occurs at more depth as compared to the ones generated by their counterparts. This further indicates that these cascades

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

Cascade Property	Cascades initiated by fringe nodes			Cascades initiated by non-fringe nodes		
	Value	σ	Percent	Value	σ	Percent
#cascades initiated	487	-	71.72%	192	-	28.28%
Avg. cascade size	177.78	1237.27	-	106.67	458.51	-
Avg. structural virality (d)	2.63	3.95	-	1.97	1.44	-
Max depth of cascade	313	-	-	84	-	-
Avg. depth of cascades	10.72	22.59	-	6.83	8.41	-
Avg. depth at which max. width was observed in cascade	5.95	14.84	-	3.02	4.61	-
Avg. #first hop nodes in cascade	5.05	8.75	-	28.05	37.29	-
Avg. #first hop non-fringe nodes	0.27	0.53	5.36%	1.44	1.59	5.14%
Avg. #first hop fringe nodes	4.77	8.79	94.64%	26.60	37.11	94.86%
Avg. #second hop nodes	8.14	22.68	-	7.87	11.08	-
Avg. #second hop nodes with non-fringe node as parent	3.64	22.03	44.77%	3.39	9.25	43.07%
Avg. #second hop nodes with fringe node as parent	4.49	7.52	55.23%	4.48	6.70	56.93%

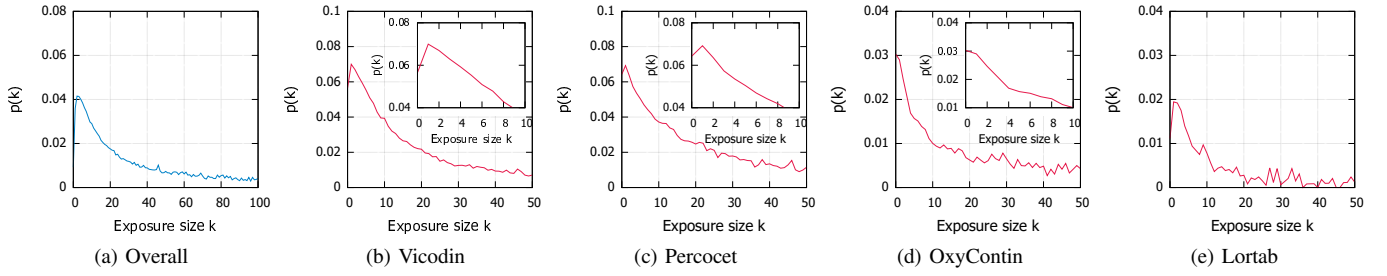


Fig. 8: Average exposure curve for drug names. $p(k)$ is the fraction of the network users who tweet about a particular drug-name directly after their k^{th} exposure to it, given that they had not tweeted about it previously. The inset in figures (b), (c) and (d) shows the behavior of structural virality near the peaks at $k = 1$.

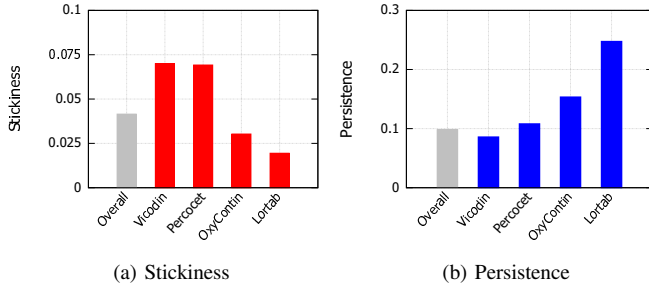


Fig. 9: Stickiness and persistence score measures for drug names as defined in [7].

survive a few initial hops with the help of other fringe nodes only to peak later with the help of certain non-fringe ones, thus revealing an organic collaboration of the non-fringe as well as fringe nodes in the spreading process.

Thus we discover that all category of users, fringe and non-fringe as well as active and non-active, contribute significantly to generating the cascades, highlighting a sizeable collective phenomenon. We next investigate the engagement behavior of users based on the contents (drug names) to discover whether Twitter serves as a more effective spreading media for certain types of drugs.

D. Stickiness and Persistence

Information diffusion concerning different content types has been studied using measures like stickiness and persistence [7], [38]. We use these measures to investigate the engagement behavior of the users with respect to the drug names. We investigate how repeated exposure to drug-abuse tweets with specific drug names influences the probability of adopting a similar engagement behavior. A user is considered to be k -exposed if there are k users, whom the current user follows, who have tweeted about drug-abuse. We use an ordinal time estimate measure for deriving the exposure curve $p(k)$, whereby, we calculate the number of users ($I(k)$) who generate their first drug-abuse tweet (an indication of adoption) after being k -exposed but before being $(k+1)$ -exposed. This value is subsequently 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 an 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.

Observations: We obtained the value of $p(k)$ for each drug type present in our dataset. Figure 8 shows the average

exposure curves for all the data and four major drug-names (determined based on #exposures). We observe in figure 9(a) Vicodin and Percocet, that were found to be mentioned in a significantly large number of tweets, have relatively much higher stickiness value (0.0699 and 0.0691, respectively) compared to the other drugs. In both the cases, peaks are found at $k = 1$, indicating that users mostly engage themselves about these tweets after a single exposure only. A similar trend was observed for OxyContin, with a peak at $k = 0$. This high value of stickiness is observed for these abused drugs due to their high popularity on Twitter. In contrast, Lortab has a relatively higher persistence of 0.24 as seen in figure 9(b), hinting that repeated exposures continue to have marginal effects on user engagement.

To explain the exceptionally high stickiness values for Vicodin (figure 8(b)) and Percocet (figure 8(c)) at $k = 1$, we looked into the tweets containing these drug names. We observed that a significantly large number of tweets mentioning Vicodin and Percocet are related to the sale of these drugs (around 40,618 and 38,910 respectively). Since these tweets are generated independently, without being exposed, we see high values of $p(k)$ at $k = 0$ and 1 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 user engagement, thus lowering the persistence. Users who are willing to discuss about these drugs rapidly engage themselves after one or two exposures. 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.

VII. CONCLUSION

This paper provides a detailed analysis of the Twitter follower network involving around 0.42 million users that are involved in the promotion of prescription drug-abuse using the Twitter platform and generating more than 50,000 cascades. We believe that this is a first major work involving such a large scale of data that details the spreading of drug-abuse messages over Twitter. Analyzing the follower network of drug-abusers reveals a heavy core structure with high local connectivity among themselves, thereby providing various alternate channel of communication among the users. Investigations on the cascades of drug-abuse tweets helped us to discover certain major findings. It was discovered that the drug-abuse tweets spread over long paths across the Twitter follower network through groups of closely connected users in the network. It was also observed that a significant percentage of cascades 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 network. The spread over those cascades has been observed to be a result of a collective phenomenon involving both the important as well as the fringe nodes, indicating a resilience to targeted elimination of few nodes. A diffusion model capturing these dynamics would

be helpful in predicting drug related cascades. Considering the limited scope of this paper, we would like to develop such models as a possible extension of the current work. Finally, observations suggest that drug-abusers on Twitter have much higher risk of adopting newer drugs as increasing exposure of them enhances the probability of adoption. These findings necessitates a deeper and more detailed study of the abuse patterns and user behavior to control the spread of this menace.

REFERENCES

- [1] D. S. Fink, J. P. Schleimer, A. Sarvet, K. K. Grover, C. Delcher, A. Castillo-Carniglia, J. H. Kim, A. E. Rivera-Aguirre, S. G. Henry, S. S. Martins *et al.*, "Association between prescription drug monitoring programs and nonfatal and fatal drug overdoses: A systematic review." *Annals of internal medicine*, vol. 168, no. 11, pp. 783–790, 2018.
- [2] 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.
- [3] 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.
- [4] 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.
- [5] 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.
- [6] M. E. Duffy and E. Thorson, "Emerging trends in the new media landscape." Springer Publishing Company New York, NY, 2009, pp. 93–116.
- [7] D. M. Romero, B. Meeder, and J. Kleinberg, "Differences in the mechanics of information diffusion across topics," in *Proc. WWW*, New York, New York, USA, 2011, p. 695.
- [8] M. D. Choudhury, M. Gamon, S. Counts, and E. Horvitz, "Predicting depression via social media." in *Proc. ICWSM*, 2013.
- [9] M. De Choudhury, "You're happy, i'm happy: Diffusion of mood expression on twitter," in *Proc HCII'15*. South Korea: Hanbit Media, Inc., 2014, pp. 169–179.
- [10] M. D. Choudhury and E. Kiciman, "The language of social support in social media and its effect on suicidal ideation risk," in *Proc. ICWSM*, 2017.
- [11] K. Saha, B. Sugar, J. Torous, B. Abraham, E. Kiciman, and M. De Choudhury, "A social media study on the effects of psychiatric medication use," in *Proc. ICWSM*, vol. 13, no. 01, 2019, pp. 440–451.
- [12] 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.
- [13] S. Rubya and S. Yarosh, "Video-Mediated Peer Support in an Online Community for Recovery from Substance Use Disorders," in *Proc. CSCW*, New York, New York, USA, 2017, pp. 1454–1469.
- [14] M. D. Choudhury and S. De, "Mental health discourse on reddit: Self-disclosure, social support, and anonymity," in *Proc. ICWSM*, 2014.
- [15] 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.
- [16] 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.
- [17] A. Sarker, K. O'Connor, 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.
- [18] C. Buntain and J. Golbeck, "This is your twitter on drugs: Any questions?" in *Proc. WWW*. ACM, 2015, pp. 777–782.
- [19] H. Hu, N. Phan, J. Geller, H. Vo, B. Manasi, X. Huang, S. Di Lorio, T. Dinh, and S. A. Chun, "Deep self-taught learning for detecting drug abuse risk behavior in tweets," in *Proc. CSoNet*. Springer, Cham, 2018, pp. 330–342.

- [20] D. Mahata, J. Friedrichs, R. R. Shah, and J. Jiang, "Detecting personal intake of medicine from twitter," *IEEE Intelligent Systems*, vol. 33, no. 4, pp. 87–95, 2018.
- [21] T. Ding, A. Roy, Z. Chen, Q. Zhu, and S. Pan, "Analyzing and retrieving illicit drug-related posts from social media," in *BIBM*. IEEE, 2016, pp. 1555–1560.
- [22] S. Chancellor, G. Nitzburg, A. Hu, F. Zampieri, and M. De Choudhury, "Discovering alternative treatments for opioid use recovery using social media," in *Proc. CHI*. ACM, 2019, p. 124.
- [23] S. Wu, J. M. Hofman, W. A. Mason, and D. J. Watts, "Who says what to whom on twitter," in *Proc. WWW*, New York, New York, USA, 2011, p. 705.
- [24] S. Aral and D. Walker, "Identifying Influential and Susceptible Members of Social Networks," *Science*, vol. 337, no. 6092, pp. 337–341, jul 2012.
- [25] E. Bakshy, J. M. Hofman, W. A. Mason, and D. J. Watts, "Everyone's an influencer," in *Proc. WSDM*. New York, New York, USA: ACM Press, 2011, p. 65.
- [26] E. Bakshy, I. Rosenn, C. Marlow, and L. Adamic, "The role of social networks in information diffusion," in *Proc. WWW*. New York, New York, USA: ACM Press, 2012, p. 519.
- [27] J. Cheng, L. Adamic, P. A. Dow, J. M. Kleinberg, and J. Leskovec, "Can cascades be predicted?" in *Proc. WWW*. New York, New York, USA: ACM Press, 2014, pp. 925–936.
- [28] 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.
- [29] D. M. Romero, W. Galuba, S. Asur, and B. A. Huberman, "Influence and passivity in social media," in *ECML/PKDD*, 2011, pp. 18–33.
- [30] D. Saez-Trumper, G. Comarela, V. Almeida, R. Baeza-Yates, and F. Benevenuto, "Finding trendsetters in information networks," in *Proc. SIGKDD*. New York, New York, USA: ACM Press, 2012, p. 1014.
- [31] M. Gomez Rodriguez, K. Gummadi, and B. Schölkopf, "Quantifying information overload in social media and its impact on social contagions," in *Proc. ICWSM*, 2014, pp. 170–179.
- [32] G. Szabo and B. A. Huberman, "Predicting the popularity of online content," *Communications of the ACM*, vol. 53, no. 8, p. 80, aug 2010.
- [33] C. Tan, L. Lee, and B. Pang, "The effect of wording on message propagation: Topic- and author-controlled natural experiments on twitter," in *Proc. ACL*. Association for Computational Linguistics, 2014, pp. 175–185.
- [34] J. Cheng, J. Kleinberg, J. Leskovec, D. Liben-Nowell, K. Subbian, L. Adamic *et al.*, "Do diffusion protocols govern cascade growth?" in *Proc. WWW*, 2018, pp. 32–41.
- [35] T.-A. Hoang and E.-P. Lim, "Tracking Virality and Susceptibility in Social Media," in *Proc. CIKM*, New York, New York, USA, 2016, pp. 1059–1068.
- [36] 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," in *PAKDD*. Springer, Cham, 2015, pp. 108–121.
- [37] 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.
- [38] 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.
- [39] S. Lai, L. Xu, K. Liu, and J. Zhao, "Recurrent convolutional neural networks for text classification," in *AAAI*, vol. 333, 2015, pp. 2267–2273.
- [40] Y. Kim, "Convolutional neural networks for sentence classification," *arXiv preprint arXiv:1408.5882*, 2014.
- [41] A. Sarker and G. Gonzalez, "Data, tools and resources for mining social media drug chatter," *BioTxtM 2016*, p. 99, 2016.
- [42] N. Phan, S. A. Chun, M. Bhole, and J. Geller, "Enabling real-time drug abuse detection in tweets," in *Proc. ICDE*. IEEE, 2017, pp. 1510–1514.
- [43] M. Pagliardini, P. Gupta, and M. Jaggi, "Unsupervised Learning of Sentence Embeddings using Compositional n-Gram Features," in *Proc. NAACL*, 2018, pp. 528–540.
- [44] 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.
- [45] O. Varol, E. Ferrara, C. A. Davis, F. Menczer, and A. Flammini, "Online human-bot interactions: Detection, estimation, and characterization," *arXiv preprint arXiv:1703.03107*, 2017.
- [46] A. Bessi and E. Ferrara, "Social bots distort the 2016 us presidential election online discussion," *First Monday*, vol. 21, no. 11-7, 2016.
- [47] M. Hintze, "Comparing the benefits of pseudonymization and anonymization under the gdpr," *Privacy Analytics, White Paper*, 2017.
- [48] 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.
- [49] F. L. Black, "Measles endemicity in insular populations: Critical community size and its evolutionary implication," *Journal of Theoretical Biology*, vol. 11, no. 2, pp. 207 – 211, 1966.
- [50] 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.
- [51] R. Peres, "The impact of network characteristics on the diffusion of innovations," *Physica A: Statistical Mechanics and its Applications*, vol. 402, pp. 330 – 343, 2014.
- [52] H. Kwak, C. Lee, H. Park, and S. Moon, "What is twitter, a social network or a news media?" in *Proc. WWW*. ACM, 2010, pp. 591–600.
- [53] Y.-X. Zhu, X.-G. Zhang, G.-Q. Sun, M. Tang, T. Zhou, and Z.-K. Zhang, "Influence of reciprocal links in social networks," *PloS one*, vol. 9, no. 7, p. e103007, 2014.
- [54] Y. W. Jinshan Qi, Xun Liang and H. Cheng, "Discrete time information diffusion in online social networks: micro and macro perspectives," *Scientific Reports*, vol. 9, no. 11872, 2018.
- [55] 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.
- [56] J. M. Kleinberg, "Authoritative sources in a hyperlinked environment," *Journal of the ACM (JACM)*, vol. 46, no. 5, pp. 604–632, 1999.
- [57] S. Wu, C. Tan, J. M. Kleinberg, and M. W. Macy, "Does bad news go away faster?" in *Proc. ICWSM*, 2011.
- [58] S. Goel, A. Anderson, J. Hofman, and D. J. Watts, "The Structural Virality of Online Diffusion," *Management Science*, p. 150722112809007, jul 2015.