Investigate Transitions into Drug Addiction through Text Mining of Reddit Data

John Lu
Department of Computer Science,
University of California
Irvine, CA
johnl19@uci.edu

Sumati Sridhar
Department of Statistics and
Operations Research, University of
North Carolina
Chapel Hill, NC
sumati@live.unc.edu

Ritika Pandey, Mohammad Al
Hasan, George Mohler
Department of Computer Science,
Indiana University - Purdue
University
Indianapolis, IN
(ripande,alhasan,gmohler)@iupui.edu

ABSTRACT

Increasing rates of opioid drug abuse and heightened prevalence of online support communities underscore the necessity of employing data mining techniques to better understand drug addiction using these rapidly developing online resources. In this work, we obtained data from Reddit, an online collection of forums, to gather insight into drug use/misuse using text snippets from users narratives. Specifically, using users' posts, we trained a binary classifier which predicts a user's transitions from casual drug discussion forums to drug recovery forums. We also proposed a Cox regression model that outputs likelihoods of such transitions. In doing so, we found that utterances of select drugs and certain linguistic features contained in one's posts can help predict these transitions. Using unfiltered drug-related posts, our research delineates drugs that are associated with higher rates of transitions from recreational drug discussion to support/recovery discussion, offers insight into modern drug culture, and provides tools with potential applications in combating the opioid crisis.

CCS CONCEPTS

• Applied computing → Health care information systems.

KEYWORDS

Drug Addiction and Recovery, Reddit Forum, Text Mining, Cox Regression

ACM Reference Format:

John Lu, Sumati Sridhar, and Ritika Pandey, Mohammad Al Hasan, George Mohler. 2019. Investigate Transitions into Drug Addiction through Text Mining of Reddit Data. In *KDD '18: ACM SIGKDD Conference on Knowledge Discovery and Data Mining, August 04–08, 2019, Anchorage, AK*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/1122445.1122456>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGKDD '19, August 04–08, 2019, Anchorage, AK

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9999-9/18/06...\$15.00

<https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

The rate of non-medical opioid use has increased markedly since the early 2000s. While recent efforts have been made to curb over-prescribing [3, 18], morbidity and mortality rates associated with opioid misuse continue to worsen [25]. Traditionally, those suffering from addiction take to support groups, such as Alcoholics Anonymous (AA) and Narcotics Anonymous (NA), on their road to recovery. These groups, which provide an encouraging community and facilitate programs for addiction management and recovery, have shown significant promise in assisting substance abusers [10].

In addition to the aforementioned support groups, federal support is provided through organizations such as the Substance Abuse and Mental Health Services Administration (SAMHSA). Among others, SAMHSA offers programs such as Medication-assisted Treatment and Too Smart To Start - the former combines the use of medication and behavioral therapy in the treatment of substance abuse and the latter is a public education initiative which deters underage alcohol use.

More recently, communities have been established in online social media and discussion forums including Reddit, MedHelp, Twitter and Drugs-Forums.com, among others. Enacted and internal stigmatization addicts must face [2, 13], in conjunction with the convenience, privacy, and anonymity of such online communities, may be explanations for their rapid expansion. These online hubs offer havens for individuals to seek advice, extend support, and share their addiction stories without fear of recourse or judgment. A previous study suggested that these communities can be tremendous resources in the understanding, monitoring and intervention of substance abuse [27].

Earlier works underscore the utility of mining data from social media sites to understand trends of human health and drug use [4, 6, 7, 14, 17, 23]; however, most of these works do not explore the particular transition from voluntary drug use to compulsive drug use, as well as the various factors that influence such a transition. Our research leverages machine learning and data mining techniques to better understand real-time, unfiltered user data presented to us via Reddit, and in particular, focuses on the understanding of how transitions into addiction can be predicted by using drug utterances and linguistic features contained in their Reddit posts. We present two statistical models: 1) a "transition classifier" that predicts if a user, given 6 months of content history in general drug discussion forums, will post in forums dedicated to substance

Subreddit	Unique Users		User Post Volume	
	% Growth 2012 to 2017	CAGR	% Growth 2012 to 2017	CAGR
Drugs	324%	34%	400%	38%
Opiates	464%	41%	657%	50%
RedditorsInRecovery	387%	37%	452%	41%
OpiatesRecovery	156%	21%	180%	23%

Figure 1: Subreddit Growth, 2012 to 2017

recovery support in the subsequent 12 months and 2) a proportional hazards survival model which estimates a user's probability of posting in a recovery forum within the next year.

2 BACKGROUND AND RELATED WORKS

Drug Addiction and the Opioid Crisis

Drug addiction is characterized by the compulsive, continued use of a substance despite its negative effects. While drug consumption often begins as recreational, frequent use increases one's tolerance, ultimately altering brain chemistry, inducing heightened desires for drugs, and prompting involuntary and compulsive use. Opioids readily increase tolerance; with each use, one requires larger doses

to reach the same level of efficacy. Consequently, opioids are highly addictive, contributing to the high prevalence of opioid misuse [5].

Since the early 2000s, the rate of opioid use has climbed, in large part, due to over-prescribing opioid pain relievers such as Oxycodone and Hydromorphone [11]. A pattern of over-prescribing and related overdose saw its peak in 2011 [22] when the Centers for Disease Control and Prevention (CDC) and the Drug Enforcement Administration (DEA) began to address this issue by implementing efforts to educate both medical professionals and the public on appropriate opiate use [5]. These efforts included tightening prescriptions and developing new prescription opioids that are "abuse-deterrent" [5]. While these efforts were met with declines

in prescription opioid abuse, illicit opioid use—including that of Heroin and Fentanyl—continued to increase and contribute to rising opioid-related injuries and overdoses [22]. One explanation for the continued rise in illicit opioids was supposedly that increased barriers to medication prompted chronic pain patients to supplement their reduced prescription allowance via self-medication, further exacerbating opioid-use risk [22]. Since Heroin is pharmacologically similar to prescription opioids, relatively cheap, and readily available, it was an obvious replacement for those previously using prescription opioids [11]. The trends outlined here demonstrate the need for balanced prevention measures that aim to reduce opioid abuse and overdose while simultaneously maintaining access to prescription opioids and treatment programs as needed [5].

Using Social Media to Understand Drug-Related Issues

Mining social media data to study health and drug-related behavior is not a new concept. MacLean et al. [14] created a model from MedHelp forums that attempts to predict addiction relapse. Paul and Dredze used a factorial LDA topic model on forums from Drugs-Forums.com to model drug types, delivery methods, and other related aspects such as cultural and health factors in the context of recreational drug use [17]. Sarker et al. assessed the use of Twitter in analyzing patterns of drug abuse, and built a binary classifier that distinguishes whether or not a tweet contains signs of medication abuse [23]. Other studies have applied machine learning to classify users of an addiction recovery mobile application [7] and to identify distinct behavioral markers between Heroin and Amphetamine dependence [1].

In the context of Reddit, Choudhury et al. used Reddit data to explore the transition between mental health discourse and suicidal ideation, and built a classifier that distinguishes between those two states [4]. Eshleman et al. demonstrated the possibility of using a binary classifier in predicting a user's propensity to post in a recovery-related subreddit [6]. Our research expands upon this latter work by exploring various linguistic factors and drug utterances in a user's post that are predictive of transitions into substance abuse; further, we build a survival model capable of estimating the probability of such transitions, providing deeper insight into the explanatory factors involved in such shifts.

3 REDDIT DATA

Reddit.com (or Reddit) is an online collection of threads grouped by user communities known as “subreddits” with each covering a distinct topic. Reddit users, or “Redditors,” subscribe and submit content to subreddits which interest them, and have their submitted content voted and commented on by fellow Redditors. Reddit has the added appeal of anonymity, allowing users to partake in unfiltered conversations on topics of shared interest. As of April 2018, the platform has over 330 million active users with 130 thousand active subreddits [9].

Threads on Reddit are defined by a user's initial post and the subsequent comments on the post by other users. Posts typically discuss a user's own substance use/abuse while comments primarily answer questions asked in the post and/or offer support to the post author. Since our objective is to learn about transitions into

substance abuse, which requires analysis of content pertaining to a user's own situation, we restrict our analysis solely to posts.

Using the pushshift.io Reddit API¹, we pulled data from January 2012 through May 2018. This dataset consisted of 309,528 posts from 125,194 unique users and included various attributes of each post such as content, title, author, date of post, number of comments, and number of upvotes.

We focus on four major drug-related subreddits: r/Opiates, r/Drugs, r/OpiatesRecovery, r/RedditorsInRecovery. In addition to having adequate user and post volume, these subreddits host discussions on a variety of drugs. Whereas r/Opiates and r/Drugs primarily serve as forums for general drug discussion, which tend to be more casual in nature, the r/OpiatesRecovery and r/RedditorsInRecovery subreddits provide an avenue for those struggling with substance abuse and addiction to seek advice, share success and relapse stories, and support others. The measurable growth in both user base and post volume between 2012 and 2017 is exhibited in Figure 1. In the proceeding analyses, we will refer to r/Opiates and r/Drugs as “casual” subreddits and r/OpiatesRecovery and r/RedditorsInRecovery as “recovery” subreddits.

4 TRANSITION CLASSIFICATION: MODELING TRANSITIONS FROM RECREATIONAL USE TO SUBSTANCE ABUSE

We trained a binary classifier to model whether a user who posts only in casual subreddits in their first 6 months will eventually go on to post in recovery subreddits in the following 12 months.

Creating Classes

Analysis was restricted to users with at least 3 posts and who exclusively posted in casual subreddits in their first 6 months on Reddit. Of these, we found the subset of users that posted in a recovery subreddit within the next 12 months; there are 220 such users, and we label this group collectively as the CAS→RECOV class. Figure 2 shows the distribution of the number of days until the first recovery post for the CAS→RECOV group.

Figure 2: Days Until First Recovery Post in the RECOV Group.

¹<https://pushshift.io/>

There are 2,836 users with at least 18 months of casual-only posts. To maintain balance between the two classes, we randomly sample 220 users from this set to form the CAS class. We then split these 440 examples into a training set (352 users) and test set (88 users).

Feature Selection

In applying a machine learning model for predicting users' transition from CAS→RECOV, each user should be represented by a vector. For this purpose, we extract various features from a user's posts, including dense embedding vectors obtained through *Doc2Vec*, linguistic characteristics of user posts, and drug utterances in the post. For the last two kinds of features, we use Kruskal-Wallis statistical test to determine which drug utterances and linguistic measures are significantly different between the CAS and CAS→RECOV classes.

Doc2Vec Embeddings. Text clustering and language processing applications involve algorithms that generally require the input text data to be represented as a fix-length vector. Bag-of-word models and n-gram models are often utilized to generate these vector representations due to their simplicity, but both lack contextual information. We therefore use Gensim² Doc2Vec model to create 100-dimensional vector representations for each user post. The idea behind such representations was first proposed by Le and Mikolov [12], who sought to create a methodology that generates vector embeddings for texts of variable length. This unsupervised framework has been shown to outperform the bag-of-words and N-gram models in sentiment analysis and other information processing tasks.

Since predictions are based on 6 months of user posts, we aggregate the document vectors of a user's posts over this period. Let $\mathcal{D}_j$ be the set of posts made by user j in the first 6 months on Reddit (ordered by date). Then, denoting the i th post of user j by d_i^j , we use the *centroid* of the doc2vec vectors of user j 's posts as a representation of the user j .

$$C_j = \frac{1}{|\mathcal{D}_j|} \sum_{i=1}^{|\mathcal{D}_j|} d_i^j \in \mathbb{R}^{100}$$

In the above equation, d_i^j is the doc2vec based vector representation of d_i^j . Though simplistic, prior works [15, 21] have suggested the efficacy of using the centroid of document representations to capture meaningful content from a set of documents.

Linguistic Measures. The specific language and words one employs in speech and writing can reveal much about one's psychological and social states [19, 20]. Further, studies have shown socio-psychological and personality differences between drug and non-drug addicts [8]. To capture both linguistic components (e.g. fraction of pronouns, verbs, and articles among others) and psychological aspects (e.g. words associated with positive/negative emotions, anger, and family/friends among others) of user posts, we use the Linguistic Inquiry and Word Count³ (LIWC), a text analysis program that categorizes words into 93 groups which reflect different emotions, thinking styles, social concerns, and parts of speech. Though LIWC

variables capture more than purely linguistic dimensions of texts, we will use the phrases "LIWC features" and "linguistic features" interchangeably hereafter for simplicity.

Keywords. We used odds ratios – a metric used in statistics to measure the association between the presence of one property with the presence of another – to find discriminative keywords for each class such as (e.g. "experiences" for CAS and "addiction" for CAS→RECOV). More concretely, if $\mathcal{W}$ is the set of all words in our training data, then the odds ratio, $OR(c, w)$, for a word, $w \in \mathcal{W}$ and class $c \in \{CAS, CAS \rightarrow RECOV\}$ is given by,

$$OR(c, w) = \frac{\frac{freq(c, w)}{freq(\neg c, w)}}{\frac{\neg freq(c, w)}{\neg freq(\neg c, w)}} = \frac{freq(c, w) * \neg freq(\neg c, w)}{freq(\neg c, w) * \neg freq(c, w)}$$

where $freq(c, w)$ is the number of posts in class c in which the word w occurs and $\neg freq(c, w)$ is the number of posts in class c where word w does not occur. The odds ratio quantifies how strongly associated w is with the class c ; a higher odds ratio implies a stronger association of w with c .

We choose w to be a keyword for c if $OR(c, w) > 2$ and $|OR(c, w) - OR(\neg c, w)| > 2$. That is, w has a high OR with respect to one class and a substantially lower OR with respect to the other class. We list a sample of associated keywords for both classes in Table 1.

Table 1: Discriminatory keywords for CAS and CAS→RECOV class using Odds Ratio

Class	Keywords
CAS	friends, completely, lsd, trip, reddit, music, weird, dr, experiences, friend, ended, 100, mdma
CAS→RECOV	quit, addiction, clean, anymore, subs, suboxone, oxy, bupe

Figure 3: Drugs with statistically significant variation in utterances between CAS and CAS→RECOV (p-values < 0.05 using Kruskal-Wallis test).

²<https://www.radmihurek.com/gensim>

³Language Inquiry and Word Count: <http://liwc.wpengin.com/>

Drug Utterances. Drug utterances serve as an indicator of which drugs users are interested in and/or currently using. For each user we calculate the mentions of each drug (as a % of all drugs they mention). This calculation includes both formal drug names of the drug, colloquial (“street”) names, and major brand names (e.g., OxyContin® is a common brand of oxycodone and “coke” is a colloquial name for cocaine). In Figure 3, we displays four drugs that have highly significant variation in utterance between the two classes. For instance, LSD is uttered more significantly in the CAS class, on the other hand, Heroin and Buprenorphine is uttered more significantly in the CAS→RECOV class.

5 SURVIVAL ANALYSIS: PREDICTING WHEN TRANSITIONS ARE LIKELY TO OCCUR

For our task, a traditional classification model predicts whether a transition would occur from CAS to CAS→RECOV. However, for such information to be useful in real-life, it is also important to know *when* such transitions might occur. This additional information can help prioritize vulnerable individuals based on their propensity to be victim of drug addiction. In this section, we present a Cox (proportional hazards) regression survival model to solve this task. This model also enables us to investigate the effect of different variables on the predicted time until one’s first recovery post.

The general formulation of a survival model centers around a random variable, T (denoting time), the survival function,

$$S(t) = Pr(T > t)$$

and the hazard function,

$$\lambda(t) = \frac{f(t)}{S(t)} \quad [where \ f(t) = \frac{d}{dt}(1 - S(t))]$$

In our case, $S(t)$ is the probability that a user will last (“survive”) more than t days without posting in a recovery subreddit and $\lambda(t)$ is the instantaneous rate of transitions to recovery subreddits (among those users that have never posted in a recovery subreddit) at time t . Adopting terminology commonly used in survival analyses, we shall say a user “survives” if he lasts more than 12 months without a recovery post and “fails” otherwise.

Right-Censored Data

Besides providing the answer whether a user will survive beyond a given time, survival model has another crucial advantage: a survival model can utilize “censored data” effectively, whereas a traditional regression model fails to do so. Censored data (also known as right-censored data) refers to observations for which the desired event has not happened yet. For instance, say, the end of observation time for our study is $\hat{t}$. Now, if a user has never posted in recovery forum within our observation period, one does not know with certainty whether that user will not post in recovery on some day $t > \hat{t}$. This is an example of *right-censored* data instances.

Our model requires users to have at least 10 posts where the first 3 were casual and not all occurring on the same day. This criteria eliminates the case of 0-day survival times, which in real-life makes little sense. In our dataset, there are 2,367 users satisfying these restrictions, and for each user, we look only at (up to) their first 12 months of posts. 165 of these users fail within the first 12 months, while the remainder are right-censored observations.

Cox Regression

A Cox regression model is a specific type of survival model that accounts for the effects of covariates on some baseline hazard function $\lambda_0(t)$. Formally, if we let $\mathbf{x}^{(i)} \in \mathbb{R}^d$ be a column vector of features for user i , the *hazard* for user i is,

$$\lambda_i(t) = \lambda_0(t) \exp\{\mathbf{x}^\top \boldsymbol{\beta}^{(i)}\}$$

and the corresponding survival function is,

$$S_i(t) = S_0(t) \exp\{-\mathbf{x}^\top \boldsymbol{\beta}^{(i)}\}$$

where $\boldsymbol{\beta} \in \mathbb{R}^d$ is a vector of trainable parameters. The data for the model can be denoted $D_{surv} = \{(y_i, \delta_i, \mathbf{x}^{(i)}) : i = 1, 2, \dots, n\}$ where y_i is the minimum of the censoring time C_i (end of observation time) and survival time T_i and

$$\delta_i = \begin{cases} 1 & \text{if } T_i = y_i \\ 0 & \text{otherwise} \end{cases},$$

where δ_i denotes whether or not an instance is censored. We train $\boldsymbol{\beta}$ by maximizing the partial likelihood estimate (under iid assumption),

$$L(\boldsymbol{\beta} | D_{surv}) = \prod_{i=1}^n \left[\frac{\lambda_0(t) \exp\{\boldsymbol{\beta}^\top \mathbf{x}^{(i)}\}}{\sum_{j \in \Psi(y_i)} \lambda_0(t) \exp\{\boldsymbol{\beta}^\top \mathbf{x}^{(j)}\}} \right]^{\delta_i}$$

where $\Psi(t) = \{i : y_i > t\}$ is the subset of users who survive past time t .

6 RESULTS

In this section we provide experimental results for both binary classification and cox regression. Results for both models suggest that drug utterances and linguistic features can be indicative of one’s propensity to shift towards substance abuse.

We use categorical accuracy and F1 score to evaluate the transition classifier. A baseline model, which used only Doc2Vec embeddings, achieved a modest 69.3% test-set accuracy. With the inclusion of LIWC linguistic features, drug utterances, and class-specific keywords in addition to tuning model parameters through grid search, we improved accuracy by roughly 5% (Table 2).

Table 2: Transition Classifier Results Summary. Table displays test-set performance of Random Forest with 170 trees (selected using grid search and 10-fold cross validation) using different features. Model trained on users’ first 6 months of posts and predicts transitions in the subsequent 12 months. Number of train and test set examples were 352 and 88, respectively.

Model	Accuracy	F1 Score
Doc2Vec	0.693	0.682
LIWC	0.659	0.659
Doc2Vec + drugs + keywords	0.716	0.725
LIWC + drugs + keywords	0.739	0.736
Doc2Vec + LIWC + drugs + keywords	0.750	0.750

Performance of the survival model was evaluated using Concordance Index (C-Index), a standard metric often employed in survival

models [26]. Our best Cox model achieved a 0.823 C-Index on the test-set (Table 3), indicating a moderately strong model.

Table 3: Cox Model Results Summary. Train/test split of 1,775 (1665 censored) and 592 (352 censored) users, respectively. C-Index shown for models using different feature sets. The model using drug utterances, keywords, and LIWC features performed best on training set using 5-fold cross validation and gave a test-set C-Index of 0.820. Test set data consisted of 45 observed and 592 censored examples.

Model	C-Index
Doc2Vec	0.790
Doc2Vec + drugs + keywords + LIWC	0.788
Drugs + keywords + LIWC	0.820
<i>Test Set Performance</i>	<i>0.820</i>

Table 4: Top 10 Explanatory Covariates

Drug Name	C-Index	LIWC feature	C-Index
Heroin	0.748	Leisure	0.668
Buprenorphine	0.702	Period	0.646
LSD	0.687	Time	0.646
Psilocybin	0.628	Ingest	0.645
Oxycodone	0.623	Informal	0.642
Marijuana	0.621	Netspeak	0.633
Ecstasy	0.614	Focuspresent	0.630
Fentanyl	0.610	Relativ	0.627
Oxymorphone	0.608	Nonflu	0.612
Amphetamine	0.597	Money	0.610

In addition to training Cox models using different feature sets, we sought to explore the explanatory strength of individual covariates by fitting a model to each individual feature (Table 4 presents the C-statistics from this experiment). Using this approach, we found that utterances of drugs such as Buprenorphine, Heroin, and LSD, have a stronger impact on a user’s predicted survival probability relative to other drugs. This is in accordance with our earlier analysis (Figure 3). Similarly, LIWC dimensions such as “leisure,” “time,” and “focuspresent” have relatively more predictive power (Table 4). These features measure the extent to which posts contain terms related to leisure activities (e.g. “cook,” “chat,” “movie”), time-related terms (e.g. “hour,” “day,” “o'clock”), and terms that indicate a focus on the present (e.g. “today,” “is,” “now”).

Survival Predictions on the Transition Dataset

We look once again at the 220 CAS and 220 CAS→RECOV Redditors, this time through the lens of our trained survival model. We fix the time duration at 12 months and compare survival probabilities between the two groups; results are shown in Figure 4. Not surprisingly, a sizable majority of the CAS group have high probability (> 90%) of surviving past 12 months.

We then approximated the addictive potential of certain drugs by measuring the average survival probability of users who share a

Figure 4: Surviving One Year. Histograms showing the number of CAS and CAS→RECOV users predicted to survive at least a year.

common drug of choice. For example, we found that users whose top drug utterances were Ecstasy or LSD had high probability of surviving while users whose drug of choice is Heroin or Buprenorphine had a comparatively smaller chance of surviving (Table 5).

Table 5: One-Year Survival Probability by Top Drug Mention

Drug Name	Surv. Prob.	Drug Name	Surv. Prob.
Ecstasy	0.987	fentanyl	0.820
LSD	0.981	cocaine	0.774
benzodiazepines	0.877	oxycodone	0.767
marijuana	0.872	Heroin	0.502
methamphetamine	0.824	Buprenorphine	0.498

7 CASE STUDY

In this section, we illustrate one of several potential uses of our models by applying them to an actual Redditor. Figure 6 presents a profile of the user including a redacted sample of his writing, the most prevalent LIWC aspects, and most uttered drugs from his posts.

The subject is a routine user of Heroin and an active participant in r/Opiates. His top 3 drug utterances are Heroin, oxycodone, and Buprenorphine with Heroin representing approximately 35% of his total drug mentions. Referring to Table 5, one may expect a lower 12-month survival probability for this user relative to someone with a random drug composition. Furthermore, LIWC dimensions of the subject’s posts are consistent with Redditors who do not survive 1 year. For example, users who fail tend to have a lower “focuspresent” dimension, and our subject scores close to the 75th-percentile in this category.

Model Predictions of Subject

Given the subject’s background and profile, it is encouraging that our transition classifier labels him as CAS→RECOV. That is, using only the subject’s first 6 months of casual posts, the classifier predicts he will post in recovery within the subsequent 12 months. Furthermore, our Cox model predicts a less than a 18.6% chance he will survive past 1 year (Figure 5), indicative of a high-risk user. Consulting the subject’s entire post history, we found that he did eventually post in r/OpiatesRecovery 200 days after his first post.

Initial Posts
Finished my stash off this morning as was expecting to have more by now. Just had a call..... Its not good. Hopefully tomorrow he said. HOPEFULLY..... Tonight is going to be a very long night for me.
...Ran out of dope last night, meaning to get more this morning. My man was out so planned on getting it while at work. But, I forgot my cash. DAMN, I have a whole night shift in front of me while starting to get sick. This sucks. Hope your all having a better night than me.
...I have aquired 15 x 100mcg Fent patches of the matrix type. I figured I would try smoking a bit of one and...I now know why its so easy to OD on them....Gonna be super careful with them though, lost 2 good friends over the years through fent OD's. Will keep posting at least once a day until their gone though. Happy nods ppl, and keep safe.

LIWC	Avg. Value	LIWC	Avg. Value
Authentic	71.261	netspeak	2.176
relativ	16.059	percept	2.028
focuspresent	13.495	leisure	1.189
Period	11.902	money	0.862
WPS	10.648	family	0.629
Sixltr	9.534	nonflu	0.579
time	8.281	sad	0.344
conj	5.841	death	0.234
informal	4.130	ingest	0.215
insight	2.491	Exclam	0.000

(a) Subset of LIWC

First Recovery Post
Hi guys, I figured this would be as good a place as any to post my story, and look for answers. I have been using Heroin on and off for about some years now...I am past the physical withdrawals now, but still really struggling with the psychological side...All my friends use, hell, most people I know use. I want to stop for my wife and son above everything else, bit finding it really hard. So, what advice do you guys have. Any will be greatly appreciated. I'm guessing I just want someone to talk to as today I am really struggling with it. Thanks

(b) Top Drugs

Figure 6: Case Study Subject Profile

Figure 5: Kaplan-Meier Curve showing surviving probability vs Days for the case study subject

8 DISCUSSION

In this work, we explored transitions from recreational substance use to substance abuse using survival analysis and binary classification. In our classifier, we were able to predict whether a previously "casual" user will post in a recovery subreddit within the next 12 months. Using our survival model, we were able to uncover distinct features such as specific drug utterances and linguistic features of a given post that influence a given user's probability of transitioning from recreational substance use to substance abuse. In this section, we discuss the implications of these results in informing drug use culture and possible methods of drug classifications, as well as highlight potential applications of our models in understanding drug use from the user's perspective.

Informing Culture and Classifications

Based on our Redditors' posts, a pattern of use was apparent: drugs such as LSD, Ecstasy, marijuana, and alcohol were often used by Redditors attending parties, music festivals, or social events in general. In contrast, opioids were commonly used as a habitual, lone activity or in smaller, low-energy gatherings. This trend learned

by inspection was further supported by post-level analyses showing a co-occurrence of frequent mentions of LSD and Ecstasy and frequent use of words such as “music,” “rave,” and “party.” Paul and Dredze found similar results in their use of topic modeling to understand various factors, including the cultural factors, associated with certain drug use. For example, the “culture” words associated with Ecstasy include [“music,” “people,” “great,” “rave”], while the “culture” words associated with opioids include [“life,” “years,” “money,” “time,” “shit”] [17].

Our survival model suggests that certain drug utterances were highly explanatory in influencing the probability of substance abuse – for instance, a user with higher Heroin mentions (as a percentage of total drug mentions) is associated with a higher probability of substance abuse, while higher LSD and Ecstasy mentions are associated with lower probabilities. This particular example directly conflicts with the current national drug classifications outlined by the DEA, who classifies drugs into five “schedules” based on their dependence potential as well as the level of accepted medical use. As it currently stands, drugs including Heroin, LSD, marijuana, and Ecstasy are classified as Schedule I, while most other opioids along with cocaine are classified as Schedule II. Our models, however, suggest that the drugs of choice for individuals who are likely to transition are often not the drugs labeled as Schedule I, with the exception of Heroin. Table 5 outlines the average “survival” probabilities associated with drugs common among our users. In our model, users whose drugs of choice are Ecstasy, LSD, and marijuana have a higher predicted survival rate which may suggest that these drugs could be better classified as Schedule II or lower. Similarly, their lower associated survival probabilities may be justification for a higher classification for drugs like Buprenorphine, Oxycodone, and Cocaine.

Other studies have reached similar conclusions. The United Kingdom has similar drug classifications based on addictive potential, and places many of the Schedule I drugs as their top dangerous drugs as well. Yet Nutt and colleagues [16] found that the UK system of classification is somewhat arbitrary and not driven by scientific evidence. In their work, they reclassified the drugs listed in the UK’s guidelines using potential of harms of individual drugs, and found that the top drugs with high potential of harm include Heroin, Cocaine, Barbiturates, Street Methadone, and Ketamine. Their findings also disagree with UK’s designation of LSD and Ecstasy in the top drug class—this aligns with our results and further suggests the higher potential of harm from opioids relative to these other drugs.

Similarly, Sarvet et al [24] discuss the increasing amount of evidence in support of the therapeutic properties of medical marijuana. According to them, medical experts, including the American Medical Association, have urged the DEA to reschedule marijuana from Schedule I to Schedule II. They argue that not only would it increase access for patients who could benefit from this form of treatment, it would also enable further research and development of cannabinoid-based medicine.

Limitations

In this section, we list some limitations to our work: 1) We identify selection bias in our Redditors, as they are choosing to express their opinions on Reddit and are likely much more open about their drug

use/recovery progress. Since this may not be indicative of all drug users, our results may not be generalizable to the larger population; 2) We use participation in general drug forums and recovery forums as proxies for recreational use and substance abuse, respectively. However, since there are no rules preventing abuse-related posts in our casual subreddits, these subreddits, though predominately about casual drug discussion, do contain abuse-related posts; 3) We cannot make any clinical diagnosis of drug addiction or mental illness based solely on our analyses; 4) Although we made attempts to include all forms of drug names in our drug counts, there may be names that are not included, and thus we may not have captured all types of drug utterances; 5) Our research uncovers certain post-level characteristics that influence the likelihood of transitions into recovery subreddits but cannot explain *why* Redditors make these transitions.

Despite these limitations, our findings provide unique insight into drug use patterns from the perspective of users and uncovered features within one’s posts that may be predictive of future substance abuse. Our work demonstrates the utility of online forums and social media sources in understanding human health and activity. Finally, the computational models that we provide can be utilized into real-life applications. For example, there can be a software App for Redditors in which they input their posts and the App returns the poster’s propensity for drug abuse. Also, this App can be used as a predictive device to help counselors and psychologists in advising their patients.

9 FUTURE WORK

Further research in this area could focus more on the temporal aspect of a user’s post sequence. Similar to the work done by Maclean et al. [14], which used a Conditional Random Field (CRF) to model phases of addiction, one could construct a sequence model using our subreddits to track users’ phases of “pre-addiction” to further analyze transitions into addiction and uncover more contextual factors that influence such transitions. We hope to explore such analysis in the future.

10 CONCLUSION

In this paper, we propose methodologies for investigating transition into drug addiction by mining voluntarily generated introspective text in Reddit forums. Our study shows that a user’s transition into drug addiction can be predicted by the keywords and linguistic features of his post, and also the name of drugs that are most frequent in his posts. Our proposed prediction models not only provides a classification, but also provide a probability value over time indicating a propensity of a user to become a drug addiction victim.

11 ACKNOWLEDGEMENT

This research was performed at Indiana University-Purdue University Indianapolis (IUPUI) in the Summer of 2018, when John and Sumati were hosted at IUPUI as REU students. Funding for this research came from Dr. Mohler and Dr. Hasan’s NSF grant no. REU-1659488.

REFERENCES

- [1] Ahn, Woo-Young, Vassileva, and Jasmin. 2016/04/01. Machine-learning identifies substance-specific behavioral markers for opiate and stimulant dependence. *Drug and Alcohol Dependence* 161 (2016/04/01). <https://doi.org/10.1016/j.drugalcdep.2016.02.008>
- [2] Magdalena Berger, Todd H. Wagner, and Laurence C. Baker. 2005. Internet use and stigmatized illness. *Social Science and Medicine* 61, 8 (2005), 1821 – 1827. <https://doi.org/10.1016/j.socscimed.2005.03.025>
- [3] Richard J. Bonnie, Morgan A. Ford, and Jonathan K. Phillips. 2017. *Pain Management and the Opioid Epidemic: Balancing Societal and Individual Benefits and Risks of Prescription Opioid Use*. The National Academies Press. <https://doi.org/10.17226/24781>
- [4] De Choudhury, Munmun, Kiciman, Emre, Mark Dredze, Glen Coppersmith, and Mrinal Kumar. 2016. Discovering Shifts to Suicidal Ideation from Mental Health Content in Social Media. In *Proc. of the CHI Conference on Human Factors in Computing Systems*. 2098–2110. <https://doi.org/10.1145/2858036.2858207>
- [5] Wilson M. Compton, Christopher M. Jones, and Grant T. Baldwin. 2016. Relationship between Nonmedical Prescription-Opioid Use and Heroin Use. *New England Journal of Medicine* 374, 2 (2016), 154–163. <https://doi.org/10.1056/NEJMr1508490>
- [6] Ryan Eshleman, Deeptanshu Jha, and Rahul Singh. 2017. Identifying individuals amenable to drug recovery interventions through computational analysis of addiction content in social media. *2017 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)* (2017). <https://doi.org/10.1109/bibm.2017.8217766>
- [7] Benjamin Fischman. 2018. Data Driven Support for Substance Addiction Recovery Communities. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems (CHI EA '18)*. ACM, New York, NY, USA, SRC07:1–SRC07:6.
- [8] M. R. Gossop and S. B. Eysenck. 1980. A Further Investigation into the Personality of Drug Addicts in Treatment. *Addiction* 75, 3 (1980), 305–311. <https://doi.org/10.1111/j.1360-0443.1980.tb01384.x>
- [9] A. Hutchinson. [n. d.]. Reddit Now Has as Many Users as Twitter, and Far Higher Engagement Rates. <https://www.socialmediatoday.com/news/reddit-now-has-as-many-users-as-twitter-and-far-higher-engagement-rates/521789/>
- [10] Samantha P Wallace Kathlene Tracy. [n. d.]. Benefits of peer support groups in the treatment of addiction. *Substance Abuse and Rehabilitation* 7 ([n. d.]). <https://doi.org/10.2147/SAR.S81535>
- [11] Andrew Kolodny, David T. Courtwright, Catherine S. Hwang, Peter Kreiner, John L. Eadie, Thomas W. Clark, and G. Caleb Alexander. 2015. The Prescription Opioid and Heroin Crisis: A Public Health Approach to an Epidemic of Addiction. *Annual Review of Public Health* 36, 1 (2015), 559–574. <https://doi.org/10.1146/annurev-publhealth-031914-122957>
- [12] Quoc Le and Tomas Mikolov. 2014. Distributed Representations of Sentences and Documents. In *Proceedings of the 31st International Conference on International Conference on Machine Learning - Volume 32 (ICML '14)*. JMLR.org, II–1188–II–1196. <http://dl.acm.org/citation.cfm?id=3044805.3045025>
- [13] Jason B. Luoma, Barbara S. Kohnenber, Steven C. Hayes, Kara Bunting, and Alyssa K. Rye. 2008. Reducing self-stigma in substance abuse through acceptance and commitment therapy: Model, manual development, and pilot outcomes. *Addiction Research & Theory* 16, 2 (2008), 149–165. <https://doi.org/10.1080/16066350701850295>
- [14] Diana MacLean, Sonal Gupta, Anna Lembke, Christopher Manning, and Jeffrey Heer. 2015. Forum77: An Analysis of an Online Health Forum Dedicated to Addiction Recovery. In *Proc. of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. 1511–1526. <https://doi.org/10.1145/2675133.2675146>
- [15] Kaustubh Mani, Ishan Verma, and Lipika Dey. 2017. Multi-Document Summarization using Distributed Bag-of-Words Model. *CoRR* abs/1710.02745 (2017). arXiv:1710.02745 <http://arxiv.org/abs/1710.02745>
- [16] David Nutt, Leslie A King, William Saulsbury, and Colin Blakemore. 2007. Development of a rational scale to assess the harm of drugs of potential misuse. *The Lancet* 369, 9566 (2007), 1047 – 1053. [https://doi.org/10.1016/S0140-6736\(07\)60464-4](https://doi.org/10.1016/S0140-6736(07)60464-4)
- [17] Michael J. Paul and Mark Dredze. 2012. Experimenting with Drugs (and Topic Models): Multi-Dimensional Exploration of Recreational Drug Discussions. In *Proc. of AAAI*.
- [18] Jonathan Penm, Neil J. MacKinnon, Jill M. Boone, Antonio Ciaccia, Cameron McNamee, and Erin L. Winstanley. 2017. Strategies and policies to address the opioid epidemic: A case study of Ohio. *Journal of the American Pharmacists Association* 57, 2, Supplement (2017), S148 – S153. <https://doi.org/10.1016/j.japh.2017.01.001>
- [19] James W. Pennebaker. 2013. *The secret life of pronouns: what our words say about us*. Bloomsbury Press.
- [20] James W. Pennebaker, Matthias R. Mehl, and Kate G. Niederhoffer. 2003. Psychological Aspects of Natural Language Use: Our Words, Our Selves. *Annual Review of Psychology* 54, 1 (2003), 547–577. <https://doi.org/10.1146/annurev.psych.54.101601.145041>
- [21] Dragomir R. Radev, Hongyan Jing, Malgorzata Stys, and Daniel Tam. 2004. Centroid-based Summarization of Multiple Documents. *Inf. Process. Manage.* 40, 6 (Nov. 2004), 919–938. <https://doi.org/10.1016/j.ipm.2003.10.006>
- [22] Mark Edmund Rose. 2018. Are Prescription Opioids Driving the Opioid Crisis? Assumptions vs Facts. *Pain Medicine* 19, 4 (2018), 793–807. <https://doi.org/10.1093/pm/pnx048>
- [23] Abeed Sarker, Karen O'Connor, Rachel Ginn, Matthew Scotch, Karen Smith, Dan Malone, and Graciela Gonzalez. 2016. Social Media Mining for Toxicovigilance: Automatic Monitoring of Prescription Medication Abuse from Twitter. *Drug Safety* 39, 3 (01 Mar 2016), 231–240. <https://doi.org/10.1007/s40264-015-0379-4>
- [24] Aaron L. Sarvet, Melanie M. Wall, David S. Fink, Emily Greene, Aline Le, Anne E. Boustead, Rosalie Liccardo Pacula, Katherine M. Keyes, Magdalena Cerda, Sandro Galea, and Deborah S. Hasin. [n. d.]. Medical marijuana laws and adolescent marijuana use in the United States: a systematic review and meta-analysis. *Addiction* 113, 6 ([n. d.]), 1003–1016. <https://doi.org/10.1111/add.14136>
- [25] P Seth, L Scholl, R Rudd, and Bacon S. 2018. Overdose Deaths Involving Opioids, Cocaine, and Psychostimulant – United States, 2015–2016. 67 (2018), 349–358. <https://doi.org/10.15585/mmwr.mm6712a1>
- [26] Harald Steck, Balaji Krishnapuram, Cary Dehing-oberije, Philippe Lambin, and Vikas C Raykar. 2008. On Ranking in Survival Analysis: Bounds on the Concordance Index. In *Advances in Neural Information Processing Systems 20*. Curran Associates, Inc., 1209–1216. <http://papers.nips.cc/paper/3375-on-ranking-in-survival-analysis-bounds-on-the-concordance-index.pdf>
- [27] Kim Jung Sunny, Marsch A. Lisa, Hancock T. Jeffrey, and Das K. Amarendra. 2017. Scaling Up Research on Drug Abuse and Addiction Through Social Media Big Data. *J Med Internet Res* 19, 10 (31 Oct 2017), e353. <https://doi.org/10.2196/jmir.6426>