Forum77: An Analysis of an Online Health Forum Dedicated to Addiction Recovery

Diana MacLean[1], Sonal Gupta[1], Anna Lembke[2], Christopher Manning[1] and Jeffrey Heer[3]
[1] Computer Science Department. Stanford University. Stanford, CA 94305 USA.
[2] Psychiatry & Behavioral Sciences Department. Stanford University. Stanford, CA 94305 USA.
[3] Computer Science & Engineering Department. University of Washington. Seattle, WA 98195 USA.
{malcdi, alembke, cmanning}@stanford.edu, sonal@cs.stanford.edu, jheer@uw.edu

ABSTRACT
Prescription drug abuse is a pressing public health issue, and people who misuse prescription drugs are turning to online forums for help. Are such forums effective? We analyze the process of opioid withdrawal, recovery and relapse on Forum77, MedHelp.org's online health forum for substance abuse recovery. Applying Prochashka's Transtheoretical Model for behavior change, we develop a taxonomy describing phases of addiction expressed by Forum77 members. We examine activity and linguistic features across the phases USING, WITHDRAWING and RECOVERING. We train statistical classifiers to identify addiction phase, relapse and whether a user was RECOVERING at the time of her last post. Applying our classifiers to 2,848 users, we find that while almost 50% relapse, the prognosis for ending in RECOVERING is favorable. Supplementing our results with users' own accounts of their experiences, we discuss Forum77's efficacy and shortcomings, and implications for future technologies.

Author Keywords
Substance abuse; addiction; cessation; online health forums

ACM Classification Keywords
H.5.3 Information Interfaces and Presentation: Group and Organization Interfaces

INTRODUCTION
Drug and alcohol use disorders, in particular the escalating misuse of prescription drugs, present one of the most pressing public health issues of the day. Addiction affects 16% of Americans ages 12 or older (about 40 million people), far exceeding the number of people afflicted with heart disease (27 million), diabetes (26 million), or cancer (19 million) [4]. In 2008, more than 36,000 deaths were due to drug overdoses; of these, opioid pain reliever (OPR) overdoses accounted for more than heroin and cocaine combined [3, 56].

CSCW '15, March 14 –18 2015, Vancouver, BC, Canada

http://dx.doi.org/10.1145/2675133.2675146

One traditional and well-accepted addiction management strategy is joining and participating in a mutual help group, such as Alcoholics Anonymous (AA) or Narcotics Anonymous (NA). However, a large and increasing number of people are going online for medical information and advice [20], and those with substance abuse problems are no exception. Several of the largest online health community websites, like MedHelp (`www.medhelp.org`), PatientsLikeMe (`www.patientslikeme.org`), and Daily Strength (`www.dailystrength.org`) include forums dedicated to the topic of addiction. It is telling that Forum77, MedHelp's substance abuse and addiction community, is the largest community on MedHelp, comprising some 740,045 posts in 80,528 discussion threads authored by 51,152 unique users.

Online mutual help groups for recovery are uniquely positioned to offer novel insight into the process of prescription drug abuse cessation for two reasons. First, prescription drug abuse is severely stigmatized, even within the medical community [44, 40]. Prior work shows that people affected by stigmatized conditions are more likely to seek help from peers [10] and to seek help online [8]. Second, patients may be unwilling to discuss issues of prescription opioid misuse with their health care providers, who may be the source of the drugs [56].

Despite the prevalence of online health forums for substance use disorders, and despite the unique role that they could play in the process of cessation, such groups are understudied. Little is known about when, in the cycle of abuse, such groups are most useful to users, what they help with, and whether they aid progression towards recovery. In this paper, we make the following contributions:

A *quantified taxonomy of phases of addiction* as expressed by users on Forum77. Our taxonomy, developed in concert with an addiction specialist, is based on Prochaska's Transtheoretical Model (TTM) of behavior change [49], and serves both as a labeling rubric for mapping text to phases of addiction, as well as a quantified summary of phase-based activity on Forum77. We use the taxonomy to manually label initiating post sequences from 191 Forum77 users (2,266 posts total) with the labels USING, WITHDRAWING or RECOVERING.

An *analysis of activity and linguistic features* across the phases of addiction. We identify features that are characteristic of each phase, and leverage them to train a Conditional Random Field (CRF) model to automatically label users' phases of addiction over their tenure on Forum77. Our CRF achieves an F1-score of 67.6% against a baseline F1-score of 20%. Using CRF-labeled sequences, we are able to identify (1) whether a user relapsed at some point during their tenure, and (2) whether a user was RECOVERING at the time of her final initiating post, with F1-scores of 78% and 82%, respectively. We make our CRF classifier freely available for download at `vis.stanford.edu/projects/forum77`.

An *analysis of transition, relapse and recovery* based on the CRF-labeled phase sequences of 2,848 Forum77 users (32,345 posts). We find that overall, progressive transitions are more prevalent than regressive transitions. Moreover, despite the fact that relapse is common (almost half of users relapse at some point during their tenure), the chances of a user RECOVERING by her final post are favorable. Finally, we observe a significant correlation between high forum engagement (both frequency of participation and volume of response posts authored) during a user's phases of USING and WITHDRAWING and the probability that she is RECOVERING when she leaves Forum77.

BACKGROUND

To our knowledge, our work is the first to investigate the topic of prescription drug abuse cessation in social media. Given the secretive and stigmatized nature of this condition [44], our contribution provides a unique and often overlooked perspective on prescription drug abuse: that of patients themselves. In this section, we provide an overview of prescription drug abuse as well as the traditional, in-person mutual help groups Alcoholics Anonymous (AA) and Narcotics Anonymous (NA). Next, we discuss work evaluating the efficacy of online health communities in general. Finally, we present work which, like ours, attempts to infer a person's health state from her social media contributions.

Prescription Drug Abuse

Prescription drug abuse (or "nonmedical use") is defined as "the use of a medication without a prescription, in a way other than prescribed, or for the experience or feelings elicited" [56]. Opioid pain relievers, such as hydrocodone, oxycodone, morphine and codeine, are the most frequently abused prescription medications [5]. In 2010, some 5.1 million Americans reported misusing prescription pain relievers in the last month, followed by sedatives (2.6 million) and stimulants (1.1 million) [5].

Recent medical research argues that drug dependence is a chronic, relapsing and remitting disorder that behaves like other chronic illnesses with a behavioral component, such as Type II Diabetes Mellitus [38]. Despite this, prescription opioid abuse is a highly stigmatized condition: the viewpoint that opioid misuse is a flaw of a person's moral character, rather than a legitimate medical condition, is common [44, 40]. The stigma is compounded by the fact that the most effective treatments for opioid use disorders are methadone or buprenorphine-assisted replacement therapies [44]. Finally, as pain treatment is often the starting point of a longer addiction to prescription opioids, it is common for people with prescription drug use disorders to acquire their drug of choice via a doctor's prescription [53, 56, 34]. In a survey of 571 individuals at an opioid detoxification clinic, Sproule et al. [53] report that 37% acquired their prescription opioids via prescription, 21% on "the street", and 25% through a combination of both.

Withdrawal

Withdrawal (or detoxification) is a painful process that is frequently compared to having a bad case of flu [6, 17]. Common withdrawal symptoms include agitation, anxiety, muscle aches, insomnia, sweating, abdominal cramping, diarrhea, goose bumps, nausea and vomiting [6]. Typically, symptom onset aligns with the first missed dose in the case of a "cold turkey" approach, or within a few days of dose reduction in the case of a taper [17]. Symptom severity peaks within a few days of final exposure, and gradually reduces as the user's physical dependence on the drug weakens [17]. Withdrawal duration, which depends on biological factors, drug, dosage levels, and withdrawal method, ranges broadly from 7-10 days (cold turkey) [21] to 20-35 days (methadone-assisted taper) [17].

Research on easing the withdrawal process focuses primarily on medication-assisted detoxification overseen by a medical professional. However, Green et al. [23] showed that informing patients in full as to the nature and severity of withdrawal symptoms that they were likely to experience resulted both in lower self-reported symptom severity scores as well as an increased probability of completing the detoxification process. Patient-reported strategies for effectively completing withdrawal include distraction and avoidance, especially in the form of physical activity [21].

Self-Detoxification

Almost no research focuses on the subject of self-detoxification. We found two studies in which attendees of the same London methadone treatment facility were interviewed about prior self-detoxification attempts. In both studies, most patients had attempted self-detoxification, and many had made multiple attempts [21, 43]. The short-term success rate per episode (24 hours of abstinence) was 41% [43], while the medium-term success rate per episode (10 days abstinence) was 24% [21]. The design of these studies naturally exclude patients who successfully maintain abstinence. When asked why their attempts had failed, subjects pointed to lack of support during detoxification [21], as well as easy access to drugs and severity of withdrawal symptoms [21, 43].

Relapse & Recovery

Relapse rates for opioid use are high. Reported reuse statistics for individuals having gone through detoxification programs range from 81–91% [22, 52]. However, long-term prognoses are more favorable, with evidence suggesting that 45–51% of patients may achieve sustained abstinence, and that sustained abstinence is a gradual process [22].

"Recovery" is a hotly contested term in drug use disorder communities. Many align with the Alcoholics Anonymous viewpoint that addiction is an uncurable disease and, as such, an individual never fully "recovers" from addiction [2].

Rather, users who reach sustained sobriety are referred to as being "in recovery". In this work, we refer to users who have overcome physical withdrawal as RECOVERING.

In-Person Mutual Help Groups

Alcoholics Anonymous (AA), founded in the 1930s [1], is one of the most utilized services for substance use disorders in the world, with over 4 million members across 100 different societies [28]. It has also given rise to other peer recovery groups for addiction, like Narcotics Anonymous (NA) and Gamblers Anonymous (GA). AA and NA are almost entirely based on mutual support, even condemning the giving of medical advice as outside the expertise of the group, instead encouraging members to see a doctor if medical or psychiatric problems arise [28].

Three decades of accumulated evidence demonstrates that active participation in such groups for addiction improves outcomes [35], although success rates are ill-defined and vary across studies [7]. A high participation level in AA is reported to be one of the strongest predictors for abstinence [45, 51]. For example, Pagano et al. [45] found that users who actively helped other AA members had a relapse rate of 55%, while those who did not relapsed at a rate of 75%. Correspondingly, many of the benefits of AA are thought to stem from the social network that it provides its members, who afford each other support, role modeling and experiential advice [30]. Kelly et al. [31] find that through their interactions with other AA members, users experience increased abstinence self-efficacy, increased spirituality/religiosity and reduced negative affect. Having a sponsor (an informal mentor) is also thought to help newcomers avoid relapse [54].

Online Health Communities

Studies of cancer-oriented online support groups show that participation assists effective disease management [27, 32], as they serve as information conduits for lifestyle management and coping [32, 41] as well as arenas for personal empowerment through narrative [18, 27]. Additional user-perceived benefits for sharing health data included suggestions for symptom management and treatment, improved understanding of one's condition, and enhanced confidence in doctors' evaluations [59]. Patients report that information provided by peers is both valuable and distinct from that given by clinicians [24]. Mankoff et al. [36] find that online peer interaction facilitates a positive shift in users' perspective about their condition, and that over time participants' roles evolve towards supporting and helping others. In concert, prior work supports the notion that people struggling with addiction may benefit from online health community participation.

Inferring Health State from Social Media

The idea that social media users' health states will be somehow reflected in the content that they contribute, and that it may be possible to predict health state from these data, has captured the interest of several researchers. De Choudhury et al. [11, 12, 13] analyze how postpartum depression (PPD) might be reflected on both Twitter and Facebook. Using their findings, they leverage activity and linguistic features to build models that predict the onset of PPD from Facebook data [13]. In other social media studies, both activity features, such as social engagement and connectivity, and linguistic features, such as affect and writing style, have been shown to be useful indicators of depression [14, 26, 50, 46], neuroticism [50] and post-traumatic stress disorder [9].

A related challenge is to identify a user's current *phase* within a specific medical condition. Jha and Elhadad [29] found that a combination of linguistic and activity features are helpful for identifying cancer stages I–IV. Murnane and Counts [42] conducted an analysis of smoking cessation as reflected on Twitter. They find that linguistic features of positive and negative sentiment, as well as social interaction variables, were significant differentiators between users who relapsed and users who ceased their smoking behavior during the time of the study. Finally, Wen and Rosê use logistic regression and flexible pattern matching over posts from an online cancer community to extract pre-defined events onto a timeline [57].

THE FORUM77 DATASET

Forum77 is an online health forum dedicated to addiction recovery. It is the largest of several topic-specific, online health forums that comprise MedHelp (`www.medhelp.org`), the world's largest online health community. In keeping with most of MedHelp's forums, Forum77 is peer moderated. We acquired our Forum77 dataset through a research agreement with MedHelp, who anonymized the data prior to sharing.

Our dataset comprises roughly 7 years of discussions on Forum77 (Jan 2007 to Apr 2014). During this time 51,153 unique users authored some 740,046 posts (80,529 threads). Figure 1 shows summary plots describing Forum77 content and user volume, average user tenure and posting rates, and thread length distribution. While most users have a tenure shorter than one month, thousands of users participate for several years. Note that we have neither demographic data (age, geographic location etc.) describing Forum77 users nor page view data describing lurking (reading without posting). As such, we restrict our analyses to Forum77 discussion content.

Discussions on Forum77 follow a simple thread structure: a user starts a discussion with an *initiating post*, along with a descriptive title. Responses are displayed below the initiating post in the order in which they are received. Features for sub-discussions (nested responses) as well as for selecting a reply as the "best response" exist, but are infrequently used.

Typically, users present *their own* substance use situation (e.g., drugs used and number of days clean) in initiating posts. In contrast, in response posts users are liable to discuss a wide range of substance abuse situations. Accordingly, we restrict our analysis of users' addiction phases to initiating posts.

EXPLORING & MODELING PHASES OF ADDICTION

To systematically analyze phases of substance abuse in Forum77, we require both a valid taxonomy of phases and a rubric mapping post text to these phases. Towards this aim, we use a set of phase labels derived from the Transtheoretical Model (TTM) of behavior change.

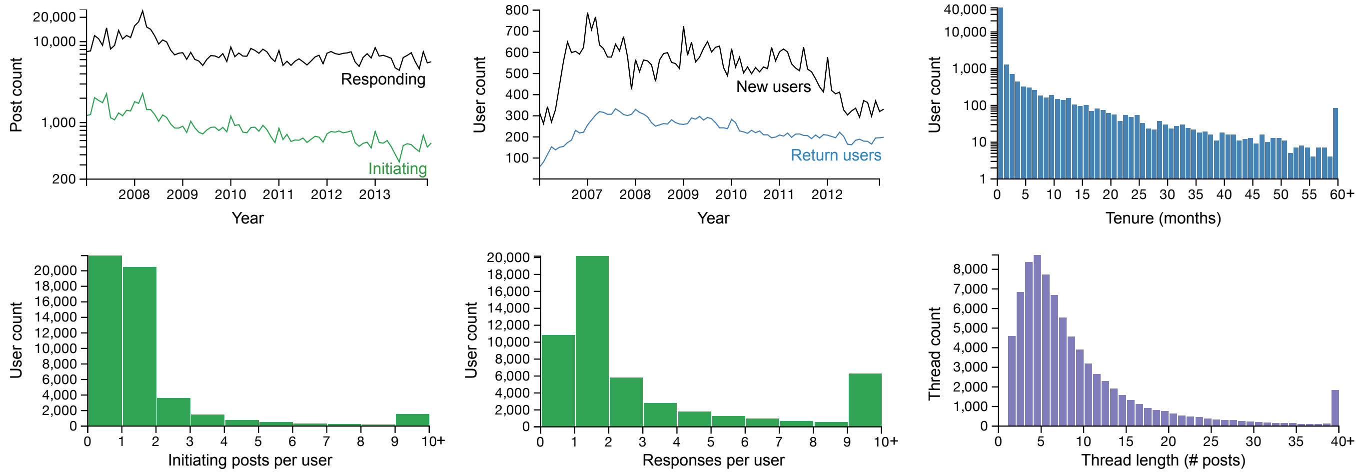

Figure 1: Summary statistics of Forum77 variables: post volume by month (top left), user volume by month (top middle), user tenure distribution (top right), user initiating post distribution (bottom left), user response post distribution (bottom middle), and thread length distribution (bottom right)

Transtheoretical Model for Behavior Change

The Transtheoretical Model (TTM) is a framework that describes six *stages of change* that a person traverses in order to manifest permanent behavior change. Established in 1997 by Prochaska & Velicer [49], the TTM has been applied to a range of behaviors, from smoking cessation [15, 42, 55] and substance abuse [39], to sustainable energy usage [25]. The intuitiveness and universal applicability of the TTM make it a useful descriptive tool; however, care should be taken before utilizing it to inform treatment or intervention [39, 58].

According to the TTM, prior to considering behavior change, a person is in the stage of **pre-contemplation**. After **contemplation**, she moves on to **preparation**, in which she makes any preparations necessary to initiate a behavior change. The person then moves on to **action**, a concerted and deliberate attempt to affect short-term behavior change. If successful, the person enters a period of **maintenance**, in which she tries to sustain the behavior change in the long term. If successful, she eventually enters the stage of **termination** [49]. As there is considerable debate over whether addiction is a terminable condition [2], we omit this stage for our purposes.

Rubric Development

In order to match Forum77 posts to TTM stages, we randomly sampled 1,000 initiating posts. Two authors mapped these posts to stages in the TTM, assigning descriptive labels to emergent sub-categories specific to the topic of addiction (e.g., *tapering* and *cold turkey* are both part of the TTM stage *Action*). We repeated this process several times, reviewing the rubric with an addiction specialist prior to finalization.

A Taxonomy of the Phases of Addiction

Table 1 describes our resulting phase taxonomy, along with example posts (synthesized from genuine posts to preserve user privacy) and the prevalence of each label in our final 1,000 initiating post sample. Although descriptively interesting, several of the labels in the taxonomy (e.g., *intent to quit* and *about to quit*) are rare. For parsimony, and to aid subsequent classification accuracy, we collapse labels into three categories: USING, WITHDRAWING and RECOVERING.

Labeling People, not Posts

Moving forward, we want to analyze addiction phases at the level of individual *people*. Two factors that emerged in our taxonomy development (see Table 1) convinced us that labeling randomly sampled posts would be insufficient for such analyses, and that we should instead label users' entire post *sequences*. The first was the high prevalence (9.8%) of *n/a* labels. These posts are often social in nature and, taken independently, are impossible to assign a class. However, when read in the context of the author's previous and following posts, the label is usually obvious. The second factor was the low prevalence of *relapse* labels. While we noticed that many users relapse, few announce the fact directly. Rather, most users will mention a relapse when they are already committed to another cessation attempt (e.g., *about to quit* or even *quitting* again). However, a relapse can still be observed in a regressive sequence (e.g., WITHDRAWING → USING). Based on these observations, in the rest of this paper we label *sequences* of posts.

CHARACTERIZING THE PHASES OF ADDICTION

Phases of addiction coincide with distinct psychological and physiological states. In this section, we analyze activity and linguistic features that might characterize an author's phase on an *initiating-day*. We define an initiating day to be any day on which the user initiated a thread on Forum77. If the author initiated multiple posts, we combine them for analysis. Our goal is two-fold: (1) to characterize phases of addiction as they are expressed on Forum77, and (2) to identify discriminative features that might be used for classification.

Sample & Labeling

To study how addiction phase changes over time, we restrict our analysis to users who have initiated at least 5 threads on

Table 1: Addiction Phase Taxonomy

Final Category	TTM phase	Label	Description	Synthesized Example	%
USING	Pre-contemplation	Using	Subject is using substances and demonstrates no intention to quit.	it has been forever since I've been here and not much has changed. I am still using the prescribed amount of oxycodone for neck pain.	3.1
		Addicted	Subject is using substances and indicates that she is addicted, but demonstrates no intent to quit.	my girlfriend and i r both addicted to percs but she is taking way more than me and keeps getting chest painonce every other week.	7.4
		Relapse	Subject has used substances again after an attempt to quit.	I just messed up majorly. I was 6 days clean, doing OK-ish, when my mother stopped by with 10 Vics "incase I needed them". Of course, being the WEAK person I am, I took them all right there.	1.3
	Contemplation	Intent to quit	Subject expresses desire to stop abusing a substance in the future.	I want off roxies. is methadone the answer. I need to work daily. I cannot do withdrawls. PLEASE HELP!	9.3
	Preparation	About to quit	Subject notes time and/or plan (e.g., tapering schedule) to quit.	i was planning to quit the first week of March. True to form addict fashion I'm out of both money and pills. So I'm about to go ct now instead of next week when I'd planned.	2.5
WITHDRAWING	Action	Quitting	Subject is in withdrawal; method unspecified.	Today is my 5th day of FREEDOM! I havent experienced any w/ds yet. So much energy.	39.1
		Tapering	Subject is in withdrawal; detoxification method is a taper.	Have some Vics I am taking. I am down to 6 a day. I plan to go down to 3 a day then 1 a day until I am done!	6.4
		Cold Turkey	Subject is in withdrawal; detoxification method is cold turkey.	I am on day 6 of CT from 150mg+ a day of ocycodone. I'm doing fine just some overall anxiousness	3.3
RECOVERING	Maintenance	In recovery	Subject has finished detoxing; no physical withdrawal symptoms expressed	Just an update to tell you that I have 67 clean days today. I feel amazing. I sleep well now and feel good! I've had a lot of discussions about aftercare.	17.8
		n/a	Impossible to determine status based on post	I've been away for few days and everything seems different. Anyway I hope everyone is doing great.	9.8

Forum77 (n=2,848 out of 29,196 users who initiated at least one post). Of these, we randomly sampled 200 users ($\sim$7% of the full 2,848) and all of their initiating posts. We discarded 9 users from the sample: two who had authored more than 100 posts, one account that belonged to MedHelp, and six accounts for which there was no clear ownership (several different people appeared to be using the same MedHelp account). The resulting sample contains 2,266 initiating posts (average 11.9 posts per user) and comprises $\sim$5.5% of the full 41,387 initiating posts authored by the 2,848 users.

Two authors categorized each initiating post in the sample using the taxonomy presented in Table 1. We labeled each user's data in chronological order so as to transfer context learned from prior labels. Disagreements (which were rare) were relabeled based on a consensus reached after discussion.

Activity Characteristics

We identify 15 activity characteristics that describe an initiator's *global* activity over time, her *local* activity 5 days prior to the initiating-day in question, and both initiator's and respondents' activity on the *initiating-day*. Table 3 describes each feature as well as feature distributions across each class.

Linguistic & Content Characteristics

Differences in word use and linguistic style are believed to reveal a range of information about people, from psychological state to social identity [48]. The Linguistic Inquiry and Word Count (LIWC) [47] software calculates 80 linguistic variables over text. In prior work, LIWC has been used to characterize and distinguish women suffering from Post-Partum Depression (PPD) [13], individuals at risk for depression [14] and smokers on Twitter who are at risk for relapse [42]. We calculate all 80 LIWC variables over initiating post text as well as any responses received on the initiating-day, then examine differences in these variables across the phases of substance abuse recovery. Due to the large number of LIWC variables, we include Tables 7 & 8 at the end of the paper.

In addition to the LIWC features, we calculate three additional variables over initiating post text. Users frequently mention how long they have been clean at the time of posting. We extract *days clean* automatically by using hand written patterns, such as "clean X days" and "X weeks off", where X represents a number. We convert X to days if necessary. We also use a more relaxed version of this feature, called *days mentioned*, in which we do not require the user to explicitly mention terms like "clean" or "off". Finally, we count the number of questions asked by identifying sentences that start with a question word and/or end with a question mark. This feature has proved helpful in prior work [13]. In subsequent experiences, we find that including these features improves classifier performance by $\sim$2.2%.

Finally, we count how many phase-specific words occur in both initiating post text as well as response text. To deter-

Table 2: Example terms typical of USING, WITHDRAWING and RECOVERING posts and their responses.

	Initiating Post	Responses
U	withdrawls, wants, hate, addicted, scared, tried, stop	situation, willing, treatment, withdrawl, option, advise, rehab, counseling
W	rls, hot, restless, aches, slept, arms, legs, headache, wd, worst, stomach, tramadol	potassium, heating, fluids, baths, pad, showers, legs, melatonin, hot, slept, bananas
R	craving, recovery, lately, sober, fight, truly, clean, cravings, true, worth	inspiration, accomplishment, congratulations, sharing, thank, miss, proud, paws

mine whether a term t is particularly descriptive of a phase p, we calculate its frequency-based odds ratio. If $f_p(t)$ is the number of posts of phase p that contain t, then:

$$OR(t,p) = \frac{f_p(t) * f_{\bar{p}}(\bar{t})}{f_p(\bar{t}) * f_{\bar{p}}(t)}$$

The odds ratio is a measure of strength of association. We calculate the odds ratio for each term across each phase, and retain terms with an odds ratio > 2. Table 2 shows sample terms for both initiating and response posts.

Results: Activity and Linguistic Features

Our feature analysis indicates that both users' activity and users' content and linguistic characteristics differ measurably across addiction phases. We present activity features in Table 3, linguistic features over initiating posts in Table 7, and linguistic features over response posts in Table 8. Unless otherwise mentioned, we use Kruskal-Wallis tests to assess statistical significance. A non-parametric test is appropriate for data that are not expected to follow a normal distribution (such as ours), and Kruskal-Wallis tests whether any pair in a trio of distributions is significantly different.

USING

This phase is characterized by long absences from the forum and, correspondingly, low levels of recent activity. Users who are USING have, on average, been absent from forum participation in all capacities for more than twice as long as users who are WITHDRAWING or RECOVERING (40 vs. $\sim$18 *days since last activity*). A longer absence from the forum may partially explain why USING posts are, on average, longer (208 vs. $\sim$180 words): users must account for lost time and bring their audience back up to speed.

Both *days clean* and *days mentioned* vary widely in USING posts, and have surprisingly high median values. Examining the underlying data provides an explanation: users who are USING often mention how long they had been clean prior to relapse in statements such as, "I was clean for 4 months before...", or, "I would have had 717 days clean today".

Finally, USING posts offer the lowest levels of positive affect (16% less than WITHDRAWING and 32% less than RECOVERING), and the highest levels of discussion around the topic of health (16% more than WITHDRAWING and 36% more than RECOVERING); characteristics that are mirrored in responses to USING posts. The lack of positivity resonates with the fact that users who are USING have either relapsed or failed to progress towards recovery.

WITHDRAWING

In recent activity, users who are WITHDRAWING issue more initiating posts and self responses than those who are USING or RECOVERING. In addition, they have the smallest average number of *days since last initiating post* (21 vs. 31 RECOVERING and 50 USING) and *days since last self-response* (29 vs. 42 RECOVERING and 66 USING).

As we might expect, WITHDRAWING users express the lowest numbers of *days clean* and *days mentioned*. In addition there is a great deal more language about feeling, biological processes and the body (Table 7). These observations align with the fact that detoxification is an uncomfortable physical process from which people constantly seek relief [17].

Responses to WITHDRAWING posts are not particularly distinctive. Aside from expressing slightly more anxiety, and containing more content about feeling and the body, other linguistic variables tend to take on a value somewhere in between those of responses to USING and RECOVERING. This may reflect that respondents try to influence users from one side of the spectrum to the other, modifying their language according to the user's progress.

RECOVERING

These users are highly active, especially in the area of responding to *other* peoples' posts. In recent activity, they issue, on average, 15.2 responses to other peoples' threads, compared to 5.5 by users who are WITHDRAWING and 1.9 by users who are USING. Moreover, unlike WITHDRAWING and USING users, their $\frac{\text{\# initiating posts}}{\text{\# responses authored}}$ tend to be < 1.

Linguistic features also suggest that RECOVERING users tend to focus on others. The pronoun *you* is used almost 100% more while the *I* pronoun is used less, and language is more social. Moreover, users express significantly more positive affect (25% more than WITHDRAWING, 48% more than USING) and less anxiety (18% less than WITHDRAWING, 16% less than USING). The evident outward focus of initiating posts from RECOVERING users resonates with the 12th step in traditional twelve-step programs such as AA, which encourage people to strengthen their sobriety by using their experiences to help others achieve it [2].

Responses to RECOVERING posts are distinct in that they express substantially more positive affect (27% more than responses to WITHDRAWING, 57% more than responses to USING). They also tend to host a notable quantity of exclamation marks (100% more than WITHDRAWING, 350% more than USING). Inspection reveals that this is an expression of excitement and encouragement in response to good news, for example, "ooooooorrrraaaahhhhh!!!!!!!!!" and "I am so PROUD of YOU!!!!!".

Table 3: Activity and content-based features for the three classes in the labeled dataset. Statistical significance is determined using Kruskal-Wallis tests (* $p < 0.05$; ** $p < 0.005$; *** $p < 0.001$) after Bonferroni corrections to adjust for family-wise error rate across all 184 variables (includes 160 LIWC variables). Column c denotes (○) if the feature is used in our CRF classifier.

				USING				WITHDRAWING				RECOVERING			
		c	p	Mean	Med	IQR	MAD	Mean	Med	IQR	MAD	Mean	Med	IQR	MAD
Activity Characteristics															
All time	# initiating posts authored		***	8.84	5.00	10.00	5.93	8.78	5.00	8.00	4.45	20.73	14.00	22.00	13.34
	# self responses authored		***	13.93	5.00	18.00	7.41	13.80	8.00	15.00	8.90	33.26	23.00	36.25	23.72
	# responses authored		***	26.90	6.00	21.00	8.90	23.61	8.00	21.00	10.38	178.69	67.00	159.25	83.77
	$\frac{\text{\# initiating posts}}{\text{\# responses authored}}$	○	***	1.36	1.00	1.31	1.02	1.28	0.82	1.26	0.85	0.53	0.22	0.35	0.21
	Days since last init. post	○	***	50.94	5.00	24.00	5.93	21.04	2.00	5.00	1.48	31.04	4.00	12.00	4.45
	Days since last self resp.	○	***	66.34	9.00	43.50	11.86	29.94	2.00	8.00	1.48	42.05	6.00	17.00	7.41
	Days since last response	○	***	73.37	5.00	27.00	5.93	33.51	2.00	6.00	1.48	28.68	2.00	5.00	1.48
	Days since last activity		***	39.56	3.00	13.00	2.97	16.66	1.00	2.00	0.00	17.76	1.00	4.00	0.00
Last 5 days	# initiating posts authored	○	***	0.93	0.00	1.00	0.00	2.01	1.00	3.00	1.48	1.81	1.00	3.00	1.48
	# self responses authored	○	***	1.37	0.00	2.00	0.00	3.32	1.00	4.00	1.48	2.89	0.00	4.00	0.00
	# responses authored	○	***	1.87	0.00	2.00	0.00	5.48	1.00	6.00	1.48	15.20	5.00	16.00	7.41
	$\frac{\text{\# initiating posts}}{\text{\# responses authored}}$	○	***	1.02	1.00	0.00	0.00	1.06	1.00	0.58	0.64	0.58	0.33	0.87	0.42
Today	# replies received		.	5.15	4.00	5.00	2.97	5.52	4.00	5.00	2.97	6.09	4.00	6.00	4.45
	# respondants		.	3.82	3.00	3.00	2.97	4.05	3.00	3.00	2.97	4.68	3.00	4.00	2.97
	# self responses		**	1.57	1.00	2.00	1.48	1.89	1.00	3.00	1.48	1.53	1.00	2.00	1.48
Post and Response Content Characteristics															
Initiating	Days clean	○	***	421.15	14.00	175.00	17.79	47.50	5.00	7.00	4.45	125.97	45.00	74.00	43.00
	Days mentioned	○	***	52.10	10.00	38.25	11.86	19.08	5.00	7.00	4.45	57.03	27.00	48.00	28.17
	# questions	○	**	2.94	2.00	3.00	1.48	2.35	2.00	2.00	1.48	2.60	2.00	2.00	1.48
	# USING terms	○	***	0.73	0.00	1.00	0.00	0.35	0.00	1.00	0.00	0.25	0.00	0.00	0.00
	# WITHDRAWING terms	○	***	0.50	0.00	1.00	0.00	1.11	1.00	2.00	1.48	0.44	0.00	1.00	0.00
	# RECOVERING terms	○	***	0.38	0.00	1.00	0.00	0.39	0.00	1.00	0.00	0.94	1.00	1.00	1.48
Responses	# USING terms	○	**	0.31	0.00	0.00	0.00	0.19	0.00	0.00	0.00	0.18	0.00	0.00	0.00
	# WITHDRAWING terms	○	***	0.86	0.00	1.00	0.00	1.18	1.00	2.00	1.48	0.76	0.00	1.00	0.00
	# RECOVERING terms	○	***	0.53	0.00	1.00	0.00	0.53	0.00	1.00	0.00	0.78	0.00	1.00	0.00

STATISTICAL CLASSIFICATION OF ADDICTION PHASE

Informed by our feature analysis, we next train a statistical classifier to automatically annotate sequences of Forum77 initiating posts with the labels USING, WITHDRAWING or RECOVERING. Analyses of phase sequences can give insight into events such as relapse and recovery. Our classifier allows us to scale such analyses to the entire Forum77 dataset. Below, we describe our classifier and report its performance. We discuss relapse and recovery in the next section.

Model & Features

A user's path through addiction phases forms a natural sequence. A Conditional Random Field (CRF) [33] is a probabilistic graphical model that performs inference over sequences, rather than individual data points. By taking into account prior and subsequent data items in a sequence, CRFs are *context sensitive*. For example, unlike a CRF, a non-sequence-based classifier might have difficulty classifying a post like, "I've been away for a few days and everything seems different. Anyway I hope everyone is doing great...", even if it was sandwiched between two posts that were obviously USING, as the post itself contains no clues as to the user's phase.

Accordingly, we train a 3-class CRF to annotate a user's sequence of initiating-days with the labels USING, WITHDRAWING or RECOVERING. We use an adapted a version of the Stanford Named Entity Recognizer package, a trainable, Java implementation of a CRF classifier, (**`http://nlp.stanford.edu/software/CRF-NER.shtml`**) that analyzes sequences of documents (default unit of analysis is a token). Tables 3, 7 and 8 indicate the subset of features that we used for classifier training. We selected features based on apparent discriminability, and iterative evaluation through 10-fold cross validation. In order to improve robustness and model potentially non-linear responses, we binned numeric features into octiles: ranks that divide the data evenly into 8 groups. While using quartiles is arguably more common in standard practice, we found that using octiles improved classifier performance.

Table 4: CRF performance scores aggregated over 10 runs of 10-fold cross validation, with randomly shuffled input sets.

Label	Precision	Recall	F1-score	Accuracy
Combined	**68.3**	**68.0**	**67.6**	**69.8**
USING	62.4	61.7	61.4	
WITHDRAWING	70.6	71.9	70.9	
RECOVERING	72.1	71.2	70.9	
Baseline	14.0	33.0	20.0	43.0

Performance

Table 4 shows Precision, F1 and Recall scores for the CRF classifier. Our classifier achieves an F1-score of 67.6%

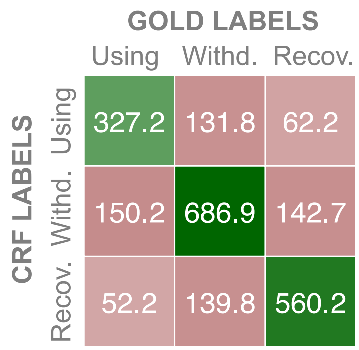

Figure 2: Confusion matrix for our CRF classifier aggregated across 10 randomized runs of 10-fold cross validation.

against a baseline F1-score of 20.0%, acquired by labeling each instance with the majority class, WITHDRAWING.

It is useful to know which labels the CRF is likely to confuse. Figure 2 shows the CRF classifier's confusion matrix. Diagonal entries indicate counts of correctly-classified instances. The strong diagonal indicates a relatively high level of accuracy. Most classification errors occur between adjacent phases: confusing USING and WITHDRAWING, and confusing WITHDRAWING and RECOVERING is common, but confusing USING and RECOVERING less so. This resonates with a point prevalent in the addiction literature: stages of recovery are not black and white but rather fall on a spectrum [16, 37].

Results

We analyze the result of applying our CRF classifier to the entirety of the Forum77 membership base who have initiated >5 posts (2,848 users, 32,345 initiating posts). Our results give us insight into common transitions between addiction phases, enabling us to answer questions such as, "If a user is WITHDRAWING today, how likely is it that she will be RECOVERING on her next initiating-day?" and "what is the most frequent phase change observed on Forum77?"

Figure 3(a) shows the normalized transition frequency matrix for USING, WITHDRAWING and RECOVERING. The most common transitions lie along the diagonal, indicating that users typically initiate consecutive posts in any one phase. Self-transitions aside, the *progressive* edges between consecutive stages (USING $\rightarrow$ WITHDRAWING and WITHDRAWING $\rightarrow$ RECOVERING) are the most common, accounting for approximately 6% and 5.2% of total transitions, respectively. In contrast, *regressive* edges between consecutive stages (WITHDRAWING $\rightarrow$ USING and RECOVERING $\rightarrow$ WITHDRAWING) are less common, accounting for 2.6% and 1.1% of total transitions, respectively.

Figure 3(b) shows transition probabilities across states. The likelihood of a same-state transition increases with the progressiveness of the state. For example, there is a 91% chance that a RECOVERING user will be RECOVERING in their next initiating post, but a 71% chance that a user in a USING state will be USING in their next initiating post.

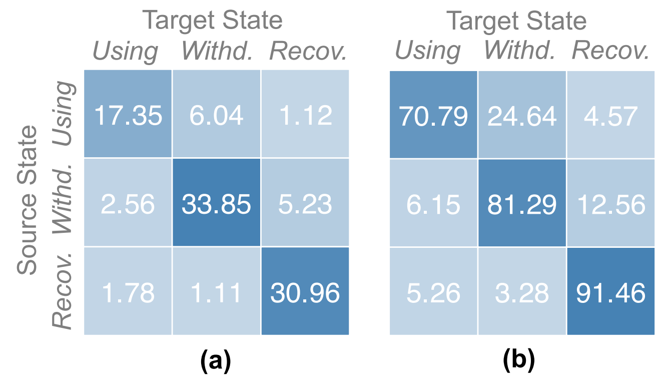

Figure 3: **(a)** Normalized transition frequencies between addiction phases (e.g., RECOVERING $\rightarrow$ USING edges comprise 1.78% of the total transitions in the CRF-labeled data) and **(b)** conditional transition probabilities (e.g., the probability of a user moving from USING to RECOVERING is 4.57%.)

Figure 4 shows the distributions of phase length in days for each phase. We calculate phase length as the number of days between the first and last post in a contiguous sequence. The typical WITHDRAWING phase lengths align well with those reported in the literature on addiction, which suggests a 7–35 day duration depending on the detoxification method used, as well as other factors [21, 17].

CLASSIFYING RELAPSE AND RECOVERY

Relapse and recovery are critical events in the process of addiction that are often viewed as "failure" or "success". Prior work in the addiction literature suggests that recovery is a long, iterative process of which relapse is a part [22]. Leveraging our CRF classifier, we present methods for identifying (1) if a user has relapsed during her tenure on the forum, and (2) if a user is RECOVERING on her last initiating-day on Forum77. We then investigate if relapse adversely correlates with a user's chance of RECOVERING. Finally, we identify activity features during USING and WITHDRAWING phases that discriminate users who last post in a state of RECOVERING.

Identifying Relapse

To identify a relapse incident, we codify three transition patterns that relate to relapse:

RECOVERING $\rightarrow$ { **WITHDRAWING, USING** }
WITHDRAWING $\rightarrow$ **USING**
WITHDRAWING $\rightarrow$ (45+ days absent) $\rightarrow$ **WITHDRAWING**

This last pattern is based on the observation that a general window for withdrawal duration is 7 - 35 days [17, 22]. As such, if a user was absent for more than 45 days, and then returned in a state of WITHDRAWING, it is likely that they failed in their initial attempt and have restarted. While it is possible that this pattern will capture individuals on a slow taper, in our experience it is unlikely that such users would be inactive for a full 45 days.

We identify whether a user relapsed or not during her tenure on Forum77 by testing whether any of the above patterns exist in her sequence of phase transitions. To evaluate the efficacy

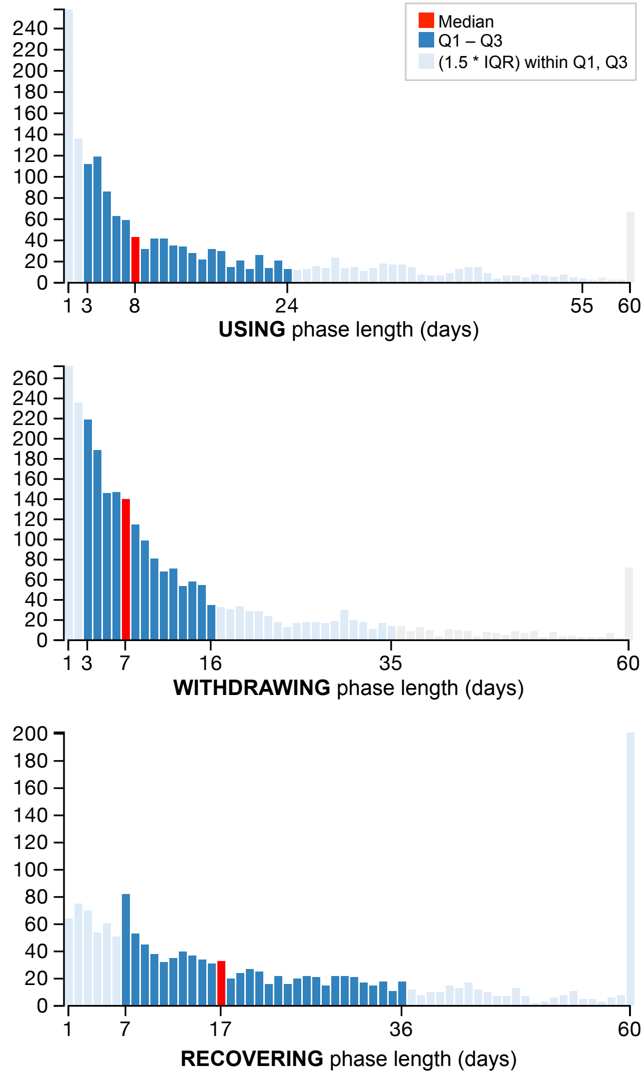

Figure 4: Distributions of phase lengths. Each red bar highlights the median value, while the dark blue region indicates the middle spread. The light blue region indicates values that fall within 1.5 $*$ the interquartile range of the middle spread.

of this approach, we apply it to both the *gold label sequences* as well as the *CRF-labeled sequences* in our labeled sample dataset. Using this technique, we achieve an F1-score of 78% and accuracy of 78% in identifying *Relapse* and *No relapse*, compared to baseline scores of 33.9% and 51.3% if we labeled each user with *No relapse*, the majority class (Table 5).

Identifying Recovery

To identify whether a user was RECOVERING when she last initiated a post on Forum77, we simply examine the final phase label in her transition sequence. Using the CRF-labeled sequences, we classify a user's last post as RECOVERING or ¬RECOVERING with an F1-score of 81.5% and accuracy of 81.6%; the comparative baselines are 34.9% and 53.4%, in which all last posts are labeled as ¬RECOVERING (Table 5).

Table 5: Performance for identifying relapse events (top) and whether a user's final state is RECOVERING (bottom).

Identifying a relapse event

Label	Precision	Recall	F1-score	Accuracy
Combined	**79.92**	**78.18**	**78.04**	**78.42**
Relapse	86.11	66.67	75.15	
No relapse	73.73	89.69	80.93	
Baseline	25.65	50.00	33.91	51.30

Identifying final initiating post phase

Label	Precision	Recall	F1-score	Accuracy
Combined	**81.47**	**81.52**	**81.49**	**81.57**
RECOVERING	79.78	80.68	80.23	
¬RECOVERING	83.17	82.35	82.76	
Baseline	26.84	50.00	34.93	53.40

Results

Using the methods described above, we identify users who are RECOVERING at the time of their last initiating post on Forum77, as well as users who have relapsed at least once during their tenure on Forum77. We apply this analysis to the entirety of the Forum77 membership base who have initiated more than 5 posts (2,848 users, 32,345 initiating posts).

Do users tend to recover on Forum77?

Overall, users progress towards recovery during their tenure. Figure 5 shows the distribution over start state, relapse, and end state for the 2,848 users described above. Most users first initiate contact on the forum when they are USING (48%), followed by WITHDRAWING (44%). In contrast, only 17% of users are USING by the time of their last post, while 37% are WITHDRAWING and 46% are RECOVERING.

Does relapsing hurt recovery likelihood?

Roughly half of users experience a relapse during their tenure. Users who experience *no relapse* are significantly more likely to end in RECOVERING than users who *relapse* (53.4% vs. 44.4% end in RECOVERING, $\chi^2_1 = 55.1, p < 0.001$). Despite this, RECOVERING is still the most likely end state for Forum77 users who relapse.

Are relapses associated with longer tenure?

Given the documented prevalence of relapse [22, 52], the observation that more than half of the users in our dataset experience *no relapse* is surprising. Analyzing tenure values reveals that the average tenure of *no relapse* users is 128 days, compared to 418 days for users who *relapse*. One hypothesis is that users who experience *no relapse* do relapse after leaving the forum and do not return.

What differentiates users who are ultimately RECOVERING?

We define a user as *active* if she initiated a post on the forum in the last 45 days of our dataset, and remove these.

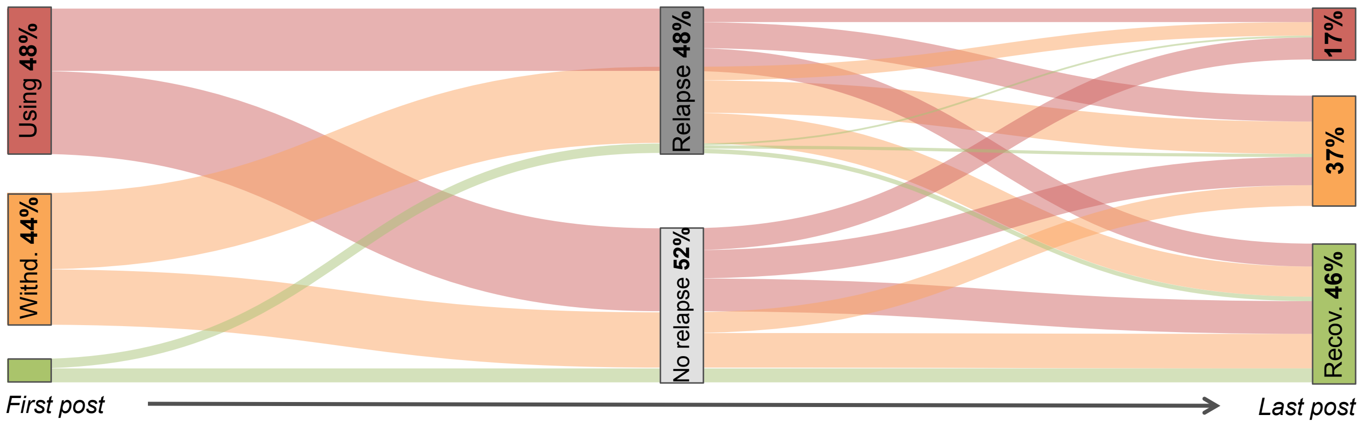

Figure 5: Aggregated user transitions from start to end state. Bar widths denote population proportion. For example, 48% of users in our sample relapsed during their tenure on Forum77.

We then analyze users' global activity characteristics (Table 3) aggregated over their USING and WITHDRAWING posts (RECOVERING posts are omitted as this is the phenomenon that we are studying). Table 6 shows the results.

Users who leave the forum in a state of RECOVERING are significantly more engaged in forum activity, even when they are USING and WITHDRAWING. The average time lapse between any form of activity (initiation, self-response and response) is about 30% shorter for those who are RECOVERING when they leave. Moreover, their activity is focused outwardly on other community members: users who are RECOVERING author, on average, 50% more responses than those who are ¬RECOVERING (average 45.6 vs. 33.8), but author slightly fewer initiating posts (average 9.0 vs. 9.9). These results resonate strongly with prior work on AA that finds that both active participation in AA and explicitly focusing on helping other members correlates with sustained abstinence [51, 45].

DISCUSSION

Use and Efficacy of Forum77

Our motivating goals were to study phases of addiction as seen on Forum77 and to analyze the forum's effectiveness in promoting recovery. In this section, we discuss Forum77's efficacy as a tool for supporting users through withdrawal, relapse and sustained recovery, drawing on post excerpts to contextualize our findings.

Supporting Withdrawal

Our results suggest that Forum77 is an effective tool for helping users through opioid withdrawals and physical detoxification. In general, users *progress* more often than they *regress* (Figure 3), and these local progressions translate into a global trend of many users reaching a state of RECOVERING during their tenure. When first initiating a post, 48% of users are USING, 44% WITHDRAWING and 8% RECOVERING; in their most recent initiating post, however, only 17% of users are USING, 37% are WITHDRAWING and 46% are RECOVERING, despite the fact that almost half of the population experiences a relapse (Figure 5). If we interpret our results as a 46% success rate on users' final detoxification attempt before leaving the forum, this is an improvement over self-detoxification success rates reported in the addiction literature [21, 43]. We must be cautious here, however, as we are comparing across differing study designs.

Forum77's efficacy at supporting detoxification may be attributable, in part, to both the strong social support and the detailed information on withdrawal that members receive from each other. Both of these factors have been shown to improve withdrawal outcomes [21, 23, 43], and qualitative remarks from users suggest that Forum77 meets the mark on both. *"I have tried to cope by myself for too long. Its so hard to deal with something like addiction by your self"*, wrote one user. *"[T]here is so much support and advice on getting through this and addiction I am living proof it works!!!!!!"*, and *"i was on here once before and was able to achieve 9 months of sobriety due to the support i had here and from meetings."* remarked others. In other cases, simply discovering a supportive community might galvanize a cessation attempt: *"up until 3 weeks ago, I had no intentions of quitting, i was just looking to find some stuff on addiction...and i just happened to run across this forum..."*.

Relapse and Shame

Despite the favorable prognosis that users are more likely to reach a state of RECOVERING during their tenure (Figure 5), we do not know whether they maintain this state upon leaving. It is possible that the same strong support network that helps users through detoxification deters them from wanting to admit a relapse. Quantitatively, although almost half of our sample relapsed (Figure 5), we rarely observed posts in which users reported a relapse immediately after the fact (Table 1).

The hypothesis that users are too ashamed to admit relapse until they implement a renewed attempt to quit is qualitatively well supported. Statements such as *"I suck!! I am so sorry, I've been too embarrased too admit I fell off the proverbial wagon around Christmas."* are common. Others, such as *"haven't posted in a few weeks because, of course, i slipped*

Table 6: Comparison of activity features for users who are RECOVERING and ¬RECOVERING in their last initiating post. Per-user values are aggregated over USING and WITHDRAWING posts. Statistical significance was determined using Kruskal-Wallis tests (* $p < 0.05$; ** $p < 0.005$; *** $p < 0.001$) after Bonferroni corrections to adjust for family-wise error rate across all 11 variables.

		RECOVERING				¬RECOVERING			
Activity Characteristic	p	**Mean**	**Med.**	**IQR**	**MAD**	**Mean**	**Med.**	**IQR**	**MAD**
# initiating posts authored	***	8.99	5	8	4.44	9.89	6	6	2.96
# self responses authored	***	19.56	8	16	10.37	17.04	9	16	8.89
# responses authored	***	45.56	9	31	13.34	33.81	8	24	10.37
$\frac{\text{\# initiating posts}}{\text{\# responses authored}}$	***	0.73	0.50	0.76	0.44	1.04	0.67	0.83	0.49
Days since last init.	***	16.39	3.33	12.41	3.95	27.05	8.30	28.36	10.53
Days since last self-response	***	17.47	3.00	13.38	3.95	29.53	8.29	31.45	10.81
Days since last response	***	15.92	1.66	7.32	2.47	25.30	4.37	21.75	5.99
Days since last activity	***	14.11	1.80	6.09	1.90	20.94	4.80	20.09	5.79
# self responses	***	1.93	1.50	1.64	1.19	1.83	1.50	1.50	1.11
# replies received	***	5.63	5.00	3.40	2.37	5.56	4.83	3.30	2.29
# respondents	***	4.09	3.83	2.00	1.60	4.01	3.70	2.03	1.42

up and am ashamed. but now i am back on track with the sub" and *"Im in day 3 of detox, i was too embarassed to post the first 3 days..."* echo these sentiments.

Supporting Sustained Recovery

Without observing users' behavior outside the forum, we cannot quantify Forum77's effectiveness at supporting long term recovery. Qualitatively, however, some users feel that this is something that Forum77 could improve upon. One user summarizes: *"I wonder if there is not a need for a forum community for long-term support. This community is great, but is skewed towards the short-term wd symptoms and getting through the initial physical pain of wd."*. Also prevalent are observations that the forum does not sufficiently prepare users to handle post-acute withdrawal syndrome (PAWS): *"I wish people would warn others about this PAWS thing"*, wrote one user. *"i was doing so good i made it to about 100 days sober ... the PAWS really got me"*, expressed another. Moreover, users who return to Forum77 after some time may find that their support network has moved on. One user who was struggling not to relapse asked *"Where are all of the friends i made here that I no longer see?!?"*.

Other users, however, give qualitative evidence in support of Forum77's efficacy at aiding sustained recovery. *"I have not posted much lately but continue to log on and read ppls posts and I believe that is a key aspect in my recovery"*, states one user. Another wrote *"when I get a craving I come here and read, even if I read it before, it helps me think of what I went through what I'm going through and how others cope"*. We found that higher engagement, in the form of activity levels and volumes of responses contributed, correlated with the chances of a user being in a phase of RECOVERING by her final initiating post. Extending this idea, one possibility is that remaining engaged with the forum (even in the form of "lurking") after reaching a state of RECOVERING helps to prevent relapses, in a similar way that continued participation in AA correlates with longer periods of sobriety [45, 51]. A deeper analysis into the mechanisms through which Forum77 does and does not support long-term recovery is an important topic for future work.

Implications for Forum Design

Our computational tools for automatically identifying addiction phases, relapses, and whether a user's tenure ends in RECOVERING could prove valuable to communities like Forum77. One question commonly asked by users is what to expect when they quit their drug of choice, and having access to this information has been shown to improve the chances of a successful cessation attempt [23]. Using phase sequence data labeled by our CRF classifier, users could set realistic expectations by exploring patterns based on thousands of users' prior experiences. Having a realistic perspective on the process of relapse and recovery may also reduce the number of instances in which users feel too embarrassed or ashamed to return to Forum77 after relapsing. Finally, exposing such data could help people find others who exhibit similar patterns to their own. Finding "people like me" is one of the primary stated reasons for user participation in online health communities [19].

While Forum77 appears to promote detoxification effectively, we observed that users have mixed feelings about how well it supports sustained recovery. It is possible that this could be addressed via altering community dynamics. For example, as we suggested above, continued participation in Forum77 post RECOVERING might help users achieve sustained recovery. Efforts focused on decreasing user churn and increasing member retention could support this. Alternatively, in a similar vein to AA's sponsorship program, which is thought to promote sustained recovery [54], we might consider automatically matching newcomers with long-term members who would act as formal mentors (or sponsors). Finally, it is possible that the community dynamics that support detoxification are different from those that would support sustained recovery. In this case, a forward reference to a different community might help RECOVERING Forum77 users plan what to do next.

Implications for Addiction Treatment

Forum77 accrues, at scale, information that is difficult to acquire through formal medical channels. First, abusing prescription drugs usually entails deceiving one's doctor. Second, addiction research data are typically acquired at point-of-care facilities (e.g., emergency rooms) or surveys at high schools or colleges. Although the ethics and privacy of such analyses must be carefully considered, it is possible that data extracted from sites like Forum77 (e.g., CRF-based transition frequencies, recovery trends, etc.) could help medical professionals and policy makers better understand patients' experiences with drug abuse. For example, insight into the day to day difficulties of opioid-assisted withdrawal might inform policy for improving the management of this popular treatment down the road. It is also possible that research like ours could illuminate poorly understood aspects of addiction: to our knowledge, ours is the first attempt to quantify the cycle of addiction.

Limitations

One limitation of this work is the selection bias of our subjects: users who come to Forum77 are likely already open to (or at least, considering) the possibility of quitting. This problem is well known to those hoping to analyze the efficacy of Alcoholics Anonymous [7]. As such, care should be taken in applying our results to a more general population who misuse prescription medication. We cannot assume, for example, that a random sample of people who misuse prescription medication would similarly progress towards recovery if they were asked to participate in Forum77. We also cannot draw epidemiological conclusions that apply to the population as a whole from these data. However, the size of Forum77, the prevalence of the opioid epidemic, and the increasing popularity of online health communities alone make the forum worth studying.

Another limitation is the acceptable—but still improvable—accuracy of our CRF classifier. While we were able to use CRF-based sequences to identify relapse and whether a final post was RECOVERING with high accuracy, improving our underlying classifier accuracy would open up more nuanced analyses. Finally, having page view data would allow us to incorporate measures of passive participation ("lurking") into our analyses, which would add a new dimension to our study. We hope to address such opportunities in future work.

CONCLUSION

We analyze the process of opioid withdrawal, recovery and relapse on Forum77, MedHelp's Addiction and Substance Abuse community. Using Prochashka's Transtheoretical Model for behavior change, we develop a taxonomy of phases of addiction that comprises three main categories: USING, WITHDRAWING and RECOVERING. The majority of initiating posts are authored when users are WITHDRAWING. Next, we analyze linguistic and behavioral features across the USING, WITHDRAWING and RECOVERING phases. Several significant differences characterize each phase, and we leverage the results of our feature analysis to train a CRF model to automatically annotate users' phase sequences. We can identify relapse events, and whether a user was RECOVERING when she authored her final post, with high accuracy from our CRF-annotated sequences.

Applying our classifier to 2,848 users reveals that progressive transitions towards RECOVERING are much more prevalent than regressive transitions. Moreover, despite the fact that almost 50% of users relapse during their tenure, leaving Forum77 in a state of RECOVERING is the most probable outcome for all users. Finally, we find that increased participation in the community correlates with a user RECOVERING by the end of her tenure: users who are RECOVERING by their final initiating post are significantly more engaged with the community when they are USING and WITHDRAWING than users who are ¬RECOVERING by their final initiating post.

To our knowledge, ours is the first work to investigate the efficacy of online mutual help groups for prescription drug abuse. Our results, which help to illuminate a previously poorly understood resource, suggest that Forum77 is an effective detoxification aid. Based on our findings, we also highlight several ways in which Forum77 might be enhanced to better support its users, such as exposing aggregate user data describing the cycle of addiction, or matching newcomers with sponsors. Finally, as the type of information shared on Forum77 is difficult to acquire at scale through traditional channels, we note that the tools and insights presented here may be of use to the addiction research community.

ACKNOWLEDGMENTS

We thank MedHelp.org for sharing their forum data. This work was supported by NSF grant number 0964173.

REFERENCES

1. Alcoholics Anonymous: Over 70 years of growth. *Alcoholics Anonymous*. [Online: `www.aa.org/aatimeline/`, accessed 15-Sept-2013].
2. The Big Book. *Alcoholics Anonymous*. [Online: `http://www.aa.org/bigbookonline`, accessed 5/20/2014].
3. Vital signs: Overdoses of prescription opioid pain relievers united states, 1999-2008. *Center for Disease Control. Morbidity and Mortality Weekly Report.* (2011). Online: `http://www.cdc.gov/mmwr/preview/mmwrhtml/mm6043a4.htm`. Accessed 93/4/2014.
4. Addiction medicine, closing the gap between science and practice. *CASAColumbia* (2012). Online: `http://www.casacolumbia.org/download/file/fid/1177`, accessed 4/5/2014.
5. Commonly abused prescription drugs. *National Institute on Drug Abuse* (2012). [Online: `http://www.drugabuse.gov/sites/default/files/rx_drugs_placemat_508c_10052011.pdf`, accessed 5/28/2014].

6. Opiate withdrawal. *MedlinePlus - U.S. National Library of Medicine* (2012). [Online: `http://www.nlm.nih.gov/medlineplus/ency/article/000949.htm`, accessed 28-May-2014].

7. Bebbington, P. E. The efficacy of alcoholics anonymous: the elusiveness of hard data. *British Journal of Psychiatry 128*, 6 (1976), 572–580.

8. Berger, M., Wagner, T. H., and Baker, L. C. Internet use and stigmatized illness. *Social Science & Medicine 61*, 8 (2005), 1821–1827.

9. Coppersmith, G., Harman, C., and Dredze, M. Measuring post traumatic stress disorder in twitter. In *Proc. ICWSM 2014*, AAAI (2014), 579–582.

10. Davison, K. P., Pennebaker, J. W., and Dickerson, S. S. Who talks? the social psychology of illness support groups. *American Psychologist 55*, 2 (2000), 205.

11. De Choudhury, M., Counts, S., and Horvitz, E. Major life changes and behavioral markers in social media: case of childbirth. In *Proc. CSCW 2013*, ACM (2013), 1431–1442.

12. De Choudhury, M., Counts, S., and Horvitz, E. Predicting postpartum changes in emotion and behavior via social media. In *Proc. CHI 2013*, ACM (2013), 3267–3276.

13. De Choudhury, M., Counts, S., Horvitz, E. J., and Hoff, A. Characterizing and predicting postpartum depression from shared facebook data. In *Proc. CSCW 2014*, ACM (2014), 626–638.

14. De Choudhury, M., Gamon, M., Counts, S., and Horvitz, E. Predicting depression via social media. In *Proc. ICWSM 2013*, AAAI (2013), 128–137.

15. DiClemente, C. C., Prochaska, J. O., Fairhurst, S. K., Velicer, W. F., Velasquez, M. M., and Rossi, J. S. The process of smoking cessation: an analysis of precontemplation, contemplation, and preparation stages of change. *Journal of Consulting and Clinical Psychology 59*, 2 (1991), 295.

16. DuPont, R. L., McLellan, A. T., White, W. L., Merlo, L. J., and Gold, M. S. Setting the standard for recovery: Physicians' health programs. *Journal of Substance Abuse Treatment 36*, 2 (2009), 159–171.

17. Farrell, M. Opiate withdrawal. *Addiction 89*, 11 (1994), 1471–1475.

18. Finfgeld, D. L. Therapeutic groups online: the good, the bad, and the unknown. *Issues in Mental Health Nursing 21*, 3 (2000), 241–255.

19. Fox, S. Peer-to-peer health care. *Pew Internet & American Life Project* (2011). [Online: `http://pewinternet.org/Reports/2011/P2PHealthcare/Summary-of-Findings.aspx`, accessed 6-January-2014].

20. Fox, S., and Duggan, M. Health online. *Pew Internet & American Life Project* (2013). [Online: `http://pewinternet.org/Reports/2013/Health-online/Summary-of-Findings.aspx`, accessed 2-April-2013].

21. Gossop, M., Battersby, M., and Strang, J. Self-detoxification by opiate addicts. a preliminary investigation. *British Journal of Psychiatry 159*, 2 (1991), 208–212.

22. Gossop, M., Green, L., Phillips, G., and Bradley, B. Lapse, relapse and survival among opiate addicts after treatment. a prospective follow-up study. *British Journal of Psychiatry 154*, 3 (1989), 348–353.

23. Green, L., and Gossop, M. Effects of information on the opiate withdrawal syndrome. *British Journal of Addiction 83*, 3 (1988), 305–309.

24. Hartzler, A., and Pratt, W. Managing the personal side of health: how patient expertise differs from the expertise of clinicians. *Journal of Medical Internet research 13*, 3 (2011).

25. He, H. A., Greenberg, S., and Huang, E. M. One size does not fit all: applying the transtheoretical model to energy feedback technology design. In *Proc. CHI 2010*, ACM (2010), 927–936.

26. Homan, C. M., Lu, N., Tu, X., Lytle, M. C., and Silenzio, V. Social structure and depression in trevorspace. In *Proc. CSCW 2014*, ACM (2014), 615–625.

27. Høybye, M. T., Johansen, C., and Tjørnhøj-Thomsen, T. Online interaction. effects of storytelling in an internet breast cancer support group. *Psycho-Oncology 14*, 3 (2005), 211–220.

28. Humphreys, K. *Circles of recovery: Self-help organizations for addictions*. Cambridge Univ. Press, 2004.

29. Jha, M., and Elhadad, N. Cancer stage prediction based on patient online discourse. In *Proc. BIO-NLP 2010*, ACL (2010), 64–71.

30. Kaskutas, L. A., Bond, J., and Humphreys, K. Social networks as mediators of the effect of alcoholics anonymous. *Addiction 97*, 7 (2002), 891–900.

31. Kelly, J. F., Hoeppner, B., Stout, R. L., and Pagano, M. Determining the relative importance of the mechanisms of behavior change within alcoholics anonymous: a multiple mediator analysis. *Addiction 107*, 2 (2012), 289–299.

32. Klemm, P., Bunnell, D., Cullen, M., Soneji, R., Gibbons, P., and Holecek, A. Online cancer support groups: a review of the research literature. *Computers Informatics Nursing 21*, 3 (2003), 136–142.

33. Lafferty, J., McCallum, A., and Pereira, F. C. Conditional random fields: Probabilistic models for segmenting and labeling sequence data.

34. Lembke, A. Why doctors prescribe opioids to known opioid abusers. *New England Journal of Medicine 367*, 17 (2012), 1580–1581.

35. Lembke, A. Humphreys, K. *Self-Help Organizations for Substance Use Disorders*. Oxford Univ. Press, 2009.

36. Mankoff, J., Kuksenok, K., Kiesler, S., Rode, J. A., and Waldman, K. Competing online viewpoints and models of chronic illness. In *Proc. CHI 2011*, ACM (2011), 589–598.

37. McLellan, A. T. What is recovery? Revisiting the Betty Ford Institute consensus panel definition. *Journal of Social Work Practice in the Addictions* (2010), 109–113.

38. McLellan, A. T., Lewis, D. C., O'Brien, C. P., and Kleber, H. D. Drug dependence, a chronic medical illness: implications for treatment, insurance, and outcomes evaluation. *JAMA 284*, 13 (2000), 1689–1695.

39. Migneault, J. P., Migneault, J. P., Adams, T. B., Migneault, J. P., Adams, T. B., Read, J. P., Migneault, J. P., Adams, T. B., and Read, J. P. Application of the transtheoretical model to substance abuse: historical development and future directions. *Drug and alcohol review 24*, 5 (2005), 437–448.

40. Miller, N. S., Sheppard, L. M., Colenda, C. C., and Magen, J. Why physicians are unprepared to treat patients who have alcohol-and drug-related disorders. *Academic Medicine 76*, 5 (2001), 410–418.

41. Mulveen, R., and Hepworth, J. An interpretative phenomenological analysis of participation in a pro-anorexia internet site and its relationship with disordered eating. *Journal of health psychology 11*, 2 (2006), 283–296.

42. Murnane, E. L., and Counts, S. Unraveling abstinence and relapse: smoking cessation reflected in social media. In *Proc. CHI 2014*, ACM (2014), 1345–1354.

43. Noble, A., Best, D., Man, L.-H., Gossop, M., and Strang, J. Self-detoxification attempts among methadone maintenance patients: what methods and what success? *Addictive Behaviors 27*, 4 (2002), 575–584.

44. Olsen, Y., and Sharfstein, J. M. Confronting the stigma of opioid use disorderand its treatment. *JAMA 311*, 14 (2014), 1393–1394.

45. Pagano, M. E., Friend, K. B., Tonigan, J. S., and Stout, R. L. Helping other alcoholics in Alcoholics Anonymous and drinking outcomes: Findings from Project MATCH. *Journal of Studies on Alcohol 65*, 6 (2004), 766.

46. Park, S., Lee, S. W., Kwak, J., Cha, M., and Jeong, B. Activities on facebook reveal the depressive state of users. *Journal of Medical Internet Research 15*, 10 (2013).

47. Pennebaker, J. W., Francis, M. E., and Booth, R. J. Linguistic inquiry and word count: Liwc 2001. *Mahway: Lawrence Erlbaum Associates 71* (2001), 2001.

48. Pennebaker, J. W., Mehl, M. R., and Niederhoffer, K. G. Psychological aspects of natural language use: Our words, our selves. *Annual Review of Psychology 54*, 1 (2003), 547–577.

49. Prochaska, J. O., and Velicer, W. F. The transtheoretical model of health behavior change. *American Journal of Health Promotion 12*, 1 (1997), 38–48.

50. Resnik, P., Garron, A., and Resnik, R. Using topic modeling to improve prediction of neuroticism and depression. In *Proceedings of the 2013 Conference on Empirical Methods in Natural*, Association for Computational Linguistics (2013), 1348–1353.

51. Sheeren, M. The relationship between relapse and involvement in alcoholics anonymous. *Journal of Studies on Alcohol and Drugs 49*, 01 (1988), 104.

52. Smyth, B., Barry, J., Keenan, E., and Ducray, K. Lapse and relapse following inpatient treatment of opiate dependence. Online: `http://www.lenus.ie/hse/bitstream/10147/110022/1/Article5260.pdf`, accessed on 4/5/2014.

53. Sproule, B., Brands, B., Li, S., and Catz-Biro, L. Changing patterns in opioid addiction characterizing users of oxycodone and other opioids. *Canadian Family Physician 55*, 1 (2009), 68–69.

54. Tonigan, J. S., and Rice, S. L. Is it beneficial to have an alcoholics anonymous sponsor? *Psychology of Addictive Behaviors 24*, 3 (2010), 397.

55. Velicer, W. F., Prochaska, J. O., Fava, J. L., Norman, G. J., and Redding, C. A. Smoking cessation and stress management: Applications of the transtheoretical model of behavior change. *Homeostasis in Health and Disease*, 5-6 (1998), 216–233.

56. Volkow, N. D. Prescription drug abuse. *National Institute on Drug Abuse Research Report Series* (2005). Online: `http://www.drugabuse.gov/sites/default/files/rxreportfinalprint.pdf`, Accessed 6/3/2014.

57. Wen, M., and Rosé, C. P. Understanding participant behavior trajectories in online health support groups using automatic extraction methods. In *Proc. GROUP 2012*, ACM (2012), 179–188.

58. West, R. Time for a change: putting the transtheoretical (stages of change) model to rest. *Addiction 100*, 8 (2005), 1036–1039.

59. Wicks, P., Massagli, M., Frost, J., Brownstein, C., Okun, S., Vaughan, T., Bradley, R., and Heywood, J. Sharing health data for better outcomes on PatientsLikeMe. *Journal of Medical Internet Research* (2010).

Table 7: LIWC features for the three classes in the labeled dataset over initiating posts. Only statistically significant variables are shown. Statistical significance is determined using Kruskal-Wallis tests (* $p < 0.05$; ** $p < 0.005$; *** $p < 0.001$) after Bonferroni corrections to adjust for family-wise error rate across all 184 variables (includes activity features). Column c denotes ($\circ$) if the feature is used in our CRF classifier.

Initiating Post Linguistic Features

			USING			WITHDRAWING			RECOVERING		
	c	p	Mean	Median	SD	Mean	Median	SD	Mean	Median	SD
Word count	○	*	208.20	151	211.06	178.92	127.00	168.81	183.23	124.50	209.24
Dic	○	***	89.26	90.17	4.89	88.10	88.89	6.26	89.38	90.54	6.59
Numerals	○	***	1.28	0.89	1.51	1.75	1.33	1.97	1.32	0.83	2.04
Function words	○	***	60.50	60.92	5.31	58.40	59.28	6.45	59.74	60.48	7.07
Pronoun		***	18.51	18.68	4.32	16.99	17.17	4.70	17.97	18.16	5.28
Personal pronoun	○	***	12.83	13.05	3.83	11.49	11.54	4.19	11.88	11.86	4.60
Pronoun: I	○	***	9.72	9.97	3.60	9.02	9.14	3.76	7.89	8.18	4.31
Pronoun: you		***	0.98	0.41	1.70	1.02	0.13	2.04	2.05	0.99	2.89
Pronoun: he/she	○	***	1.14	0	2.08	0.74	0	1.82	1	0	2.46
Pronoun: they	○	***	0.65	0.20	1.05	0.47	0	1.12	0.54	0	1.03
Pronoun: impersonal		*	5.68	5.33	2.82	5.49	5.26	2.81	6.09	5.76	3.35
Verb	○	**	18.54	18.69	3.76	17.64	17.59	4.20	18.13	17.96	4.91
Present tense	○	***	12.56	12.55	3.90	11.53	11.24	4.09	11.95	11.63	4.45
Numbers	○	**	0.71	0.48	0.93	0.75	0.37	1.12	0.54	0	0.89
Social	○	***	7.60	6.59	4.79	6.38	5.26	5.18	8.85	7.89	5.90
Humans	○	*	0.49	0	0.76	0.40	0	0.79	0.57	0	1.04
Affect	○	***	5.30	5.00	2.76	5.76	5.54	3.09	6.41	6.11	3.52
Affect: positive	○	***	2.80	2.45	1.99	3.33	2.86	2.85	4.14	3.50	3.16
Affect: anxiety	○	**	0.61	0.25	0.88	0.55	0	0.98	0.45	0	0.90
Cognitive Mechanisms	○	*	17.27	16.98	4.50	17.14	17.09	4.95	17.93	17.96	5.11
Certain	○	*	1.21	1.03	1.22	1.41	1.21	1.41	1.57	1.36	1.53
Inhibition	○	*	0.50	0.23	0.70	0.41	0	0.74	0.43	0	0.76
See	○	*	0.34	0	0.65	0.30	0	0.80	0.50	0	1.14
Feel	○	***	0.73	0.45	1.10	1.18	0.83	1.50	0.85	0.50	1.23
Biological	○	***	3.87	3.46	2.63	4.01	3.70	2.90	3.31	2.89	2.72
Body	○	***	0.58	0	1	1.13	0.63	1.53	0.68	0	1.12
Health	○	***	3.00	2.63	2.29	2.58	2.13	2.36	2.20	1.72	2.25
Relative	○	***	13.46	13.39	4.65	15.04	14.75	5.25	13.72	13.61	5.23
Time	○	***	7.24	6.86	3.46	8.51	7.87	4.21	7.33	7.02	4.23
Home	○	***	0.30	0	0.54	0.40	0	0.77	0.68	0.14	1.18
Comma	○	**	3.01	2.17	3.36	2.75	1.94	3.27	2.19	1.63	2.43
QMark	○	*	1.35	0.52	2.87	1.34	0.40	2.58	1.50	0	4.92
Other Punctuation	○	***	0.81	0	1.77	0.89	0	1.91	0.62	0	2.05

Table 8: LIWC features for the three classes in the labeled dataset over response posts. Only statistically significant variables are shown. Statistical significance is determined using Kruskal-Wallis tests (* $p < 0.05$; ** $p < 0.005$; *** $p < 0.001$) after Bonferroni corrections to adjust for family-wise error rate across all 184 variables (includes activity features). Column c denotes (○) if the feature is used in our CRF classifier.

Response Post Linguistic Features

			USING			WITHDRAWING			RECOVERING		
	c	p	Mean	Median	SD	Mean	Median	SD	Mean	Median	SD
Word count		***	494.69	347.00	506.67	427.38	284.00	487.46	356.29	210.50	439.75
Words per sentence		***	19.21	15.40	18.60	17.04	14.09	14.73	14.98	12.99	14.25
Numerals	○	*	0.75	0.43	1.02	0.95	0.68	1.12	0.95	0.56	1.49
Function words		***	59.01	59.85	4.56	56.95	57.69	5.41	55.82	57.06	7.17
Personal Pronouns		***	10.86	11.36	3.99	10.21	10.53	3.81	10.86	11.58	4.71
Pronoun: she/he		**	0.68	0	1.35	0.44	0	1.16	0.64	0	1.63
Pronoun: they		***	0.66	0.41	0.91	0.49	0.27	0.66	0.49	0.13	0.90
Pronoun: impersonal		**	5.48	5.67	2.20	5.57	5.78	2.36	5.10	5.32	2.75
Article		***	4.91	4.98	2.06	4.75	4.96	2.02	4.20	4.41	2.23
Verb		**	17.26	18.15	4.94	17.13	17.82	4.88	16.09	17.23	5.78
Aux. verb		***	10.67	11.11	3.51	10.37	10.68	3.44	9.66	10.33	3.96
Future		***	1.50	1.44	1.07	1.50	1.43	1.13	1.10	1.01	1.03
Preposition		***	11.63	12.27	3.57	11.19	11.66	3.38	10.61	11.51	4.14
Conjunction		***	6.39	6.76	2.33	6.18	6.58	2.46	5.72	6.13	2.69
Quantitative		***	3.00	2.99	1.52	2.94	2.88	1.64	2.50	2.58	1.67
Social	○	***	10.26	10.11	4.77	8.83	8.75	4.23	9.78	9.81	5.45
Affect	○	***	5.73	5.76	2.68	6.55	6.34	3.25	7.54	7.33	4.31
Affect: positive	○	***	3.72	3.53	2.43	4.61	4.10	3.17	5.84	5.13	4.36
Affect: negative	○	***	1.96	1.92	1.34	1.90	1.87	1.33	1.67	1.50	1.51
Affect: anxiety	○	***	0.36	0.24	0.47	0.40	0.23	0.55	0.32	0	0.61
Cognitive Processes		***	19.37	17.81	7.77	18.71	17.43	7.83	18.77	16.80	10
Discrepancy		***	2.32	2.32	1.31	1.92	1.88	1.33	1.63	1.60	1.30
Tentative		***	3.35	3.25	1.79	3.12	3.09	1.77	2.55	2.45	1.96
Exclusive		***	3.35	3.40	1.62	3.07	3.18	1.66	2.56	2.60	1.83
Perceptual processes		***	1.52	1.48	1.07	1.90	1.81	1.34	1.87	1.68	1.55
Feel		***	0.64	0.53	0.70	0.91	0.76	0.85	0.65	0.45	0.76
Biological		***	3.46	3.20	2.17	3.42	3.22	2.46	2.71	2.41	2.39
Body		***	0.52	0.28	0.78	0.78	0.45	1.08	0.52	0.19	0.90
Health	○	***	2.68	2.45	1.85	2.24	1.95	1.90	1.70	1.32	1.76
Sexual		***	0.15	0	0.35	0.14	0	0.36	0.30	0	0.89
Ingetion		*	0.17	0	0.39	0.30	0	0.66	0.25	0	0.71
Relativity		**	11.46	11.82	4.39	12.36	12.68	4.70	11.90	12.50	5.37
Time		**	5.29	5.10	2.90	5.88	6.06	3.12	5.66	5.69	3.33
Money		*	0.32	0.13	0.55	0.28	0	0.56	0.23	0	0.42
Assent	○	***	0.27	0.07	0.50	0.40	0.18	0.81	0.62	0.27	2.01
Colon		**	0.09	0	0.20	0.15	0	0.42	0.27	0	0.84
Exclamation	○	***	1.02	0.34	1.79	2.25	0.82	5.08	4.52	1.68	8.40
Dash		**	0.79	0.28	2.08	0.82	0	2.20	0.62	0	1.64
Other punctuation	○	***	3.41	2.84	2.55	4.29	3.53	3.22	5.64	4.29	6.35
All punctuation	○	***	22.07	21.51	9.71	25.75	23.69	14.52	29.69	26.82	19.27